An Unsteady Continuous Adjoint Approach for Aerodynamic Design

Thomas D. Economon

NASA AMS Seminar
June 9th, 2015

Outline

- Design in Unsteady Flows
- Mathematical Formulation
- Numerical Results
- Future Directions for Adjoints

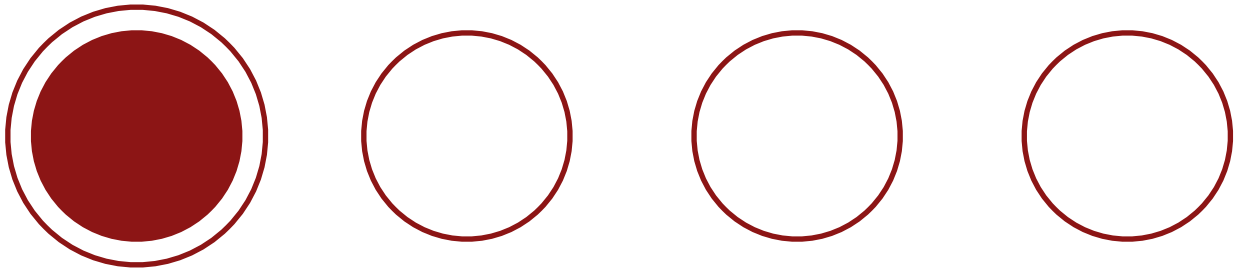

Design in Unsteady Flows

Design in Unsteady Flows

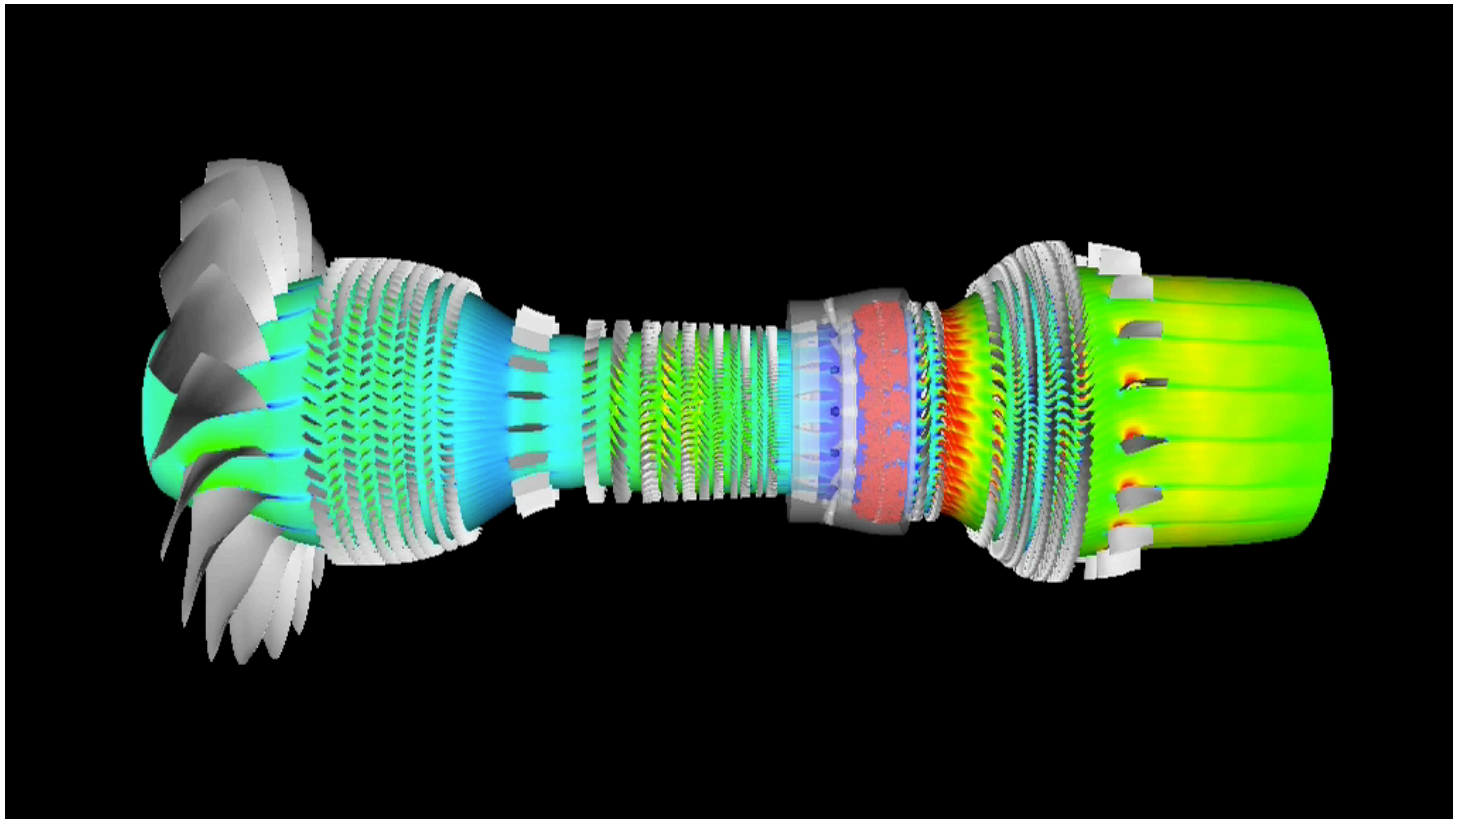

- There is a growing interest for design in unsteady flows, and it is becoming more tractable with increases in computing power
- To achieve higher efficiencies, many critical applications could immediately benefit from a time-accurate design approach: turbomachinery, open rotors, rotorcraft, wind turbines, maneuvering flight, flapping flight, etc.
- Unsteady treatment will also directly enable multidisciplinary design, analysis, and optimization (MDO) involving other time-dependent physics (structures, acoustics, flow control techniques)

Animation

Optimal Shape Design

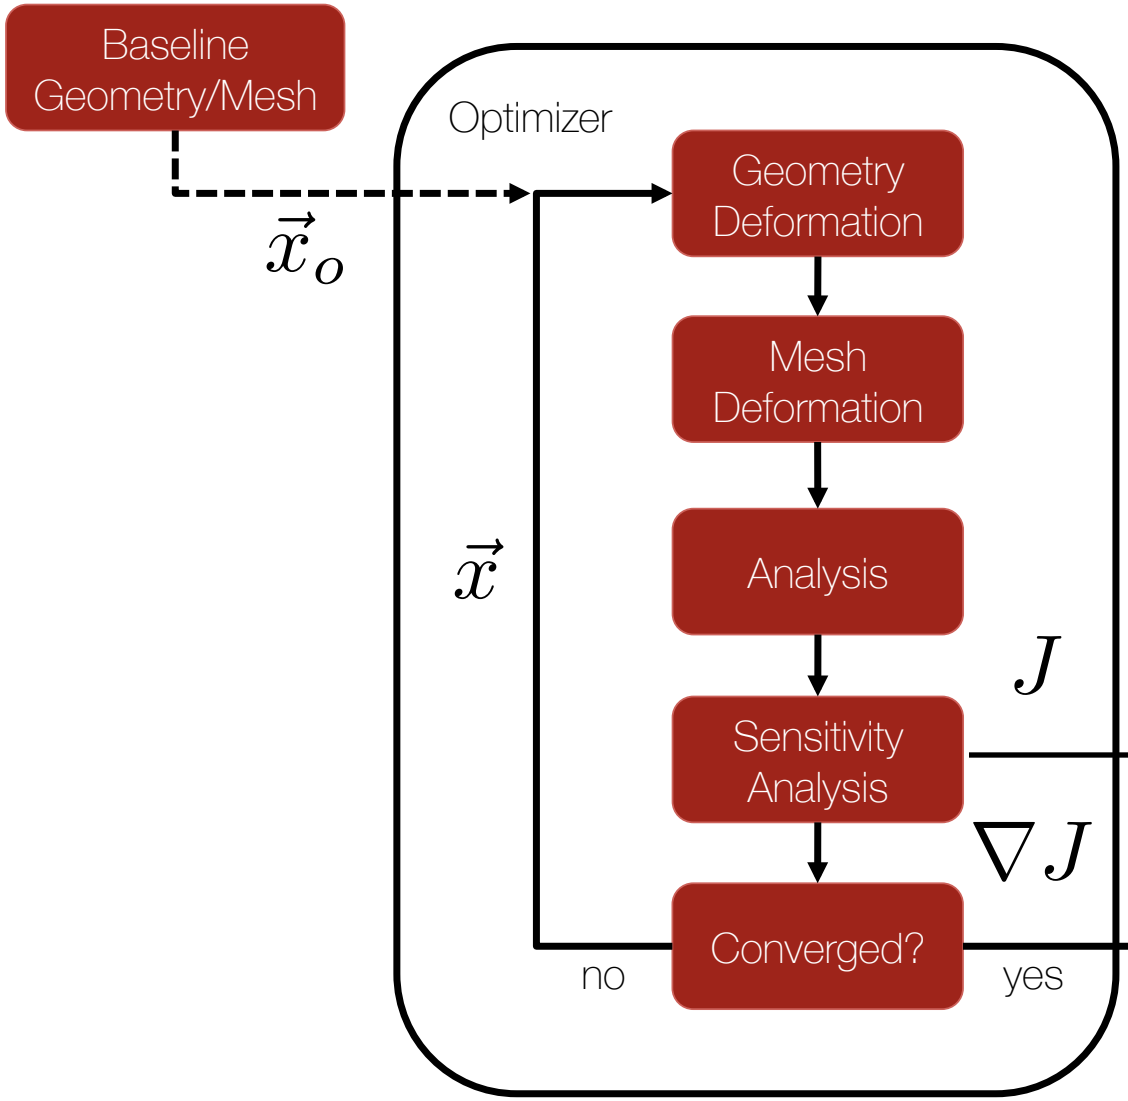

$\vec{x}$ vector of design variables

J objective/constraint function(s)

CFD gives us J (lift, drag, etc.), but how do we get ∇J efficiently?

The Adjoint Approach!

The Adjoint Approach

- Features of the **continuous adjoint** (linearize then discretize):
  - treat the continuous PDEs for the flow equations and recover a PDE system for the adjoint by using variational methods
  - obtain analytic expression for the gradient as a surface integral, i.e., a **surface formulation**
  - flexibility in choice of numerical methods
  - tedious derivation, especially boundary conditions
- Features of the **discrete adjoint** (discretize then linearize):
  - treat the discretized form of the flow equations and recover a linear system for the adjoint
  - yields numerically exact gradient of the discretized functional
  - can benefit from the use of algorithmic differentiation (AD)
  - can suffer from large memory & compute overhead (large, possibly stiff linear system)

Design in Unsteady Flows

- Challenges for unsteady design:
  - Computational cost can increase dramatically for time-accurate simulations.
  - Need to manage large amounts of solution data (adjoint requires reverse time integration).
  - Handling moving surfaces / dynamic meshes in the formulation requires the **Arbitrary Lagrangian-Eulerian (ALE)** form of the flow equations.

Design in Unsteady Flows

- For large-scale problems, need to avoid computational overhead and memory bottlenecks
- Due to the complexity of the unsteady design problem, a continuous adjoint approach is appealing due to…
  - recovering a surface formulation for the gradient with no dependence on the volume mesh (use shape calculus)
  - flexibility in numerical methods to help mitigate convergence issues for stiff problems
  - the time-accurate, continuous adjoint PDE can also be discretized for different problems immediately (non-inertial or time-spectral approaches, for instance)

Design in Unsteady Flows

- Unsteady Adjoint State of the Art:
  - Nadarajah and Jameson (2007) performed shape design for pitching airfoils using both continuous and discrete adjoints
  - Rumpfkiel and Zingg (2007) used a discrete adjoint for the control of unsteady flows in 2-D, including drag minimization and noise minimization for airfoils
  - Mani and Mavriplis (2008) and Mavriplis (2008) showed unsteady discrete adjoints for two- and three-dimensional, deforming, unstructured meshes (pitching wing)
  - Nielsen et al. (2010, 2012) developed a discrete adjoint approach for turbulent flows on dynamic, possibly overset, deforming meshes
  - **Economon et al. (2013)** investigated an unsteady continuous adjoint (inviscid) for the design of pitching airfoils in the presence of sliding mesh interfaces
  - **Economon et al. (2014)** developed a continuous adjoint surface formulation for the unsteady, compressible RANS equations on dynamic meshes

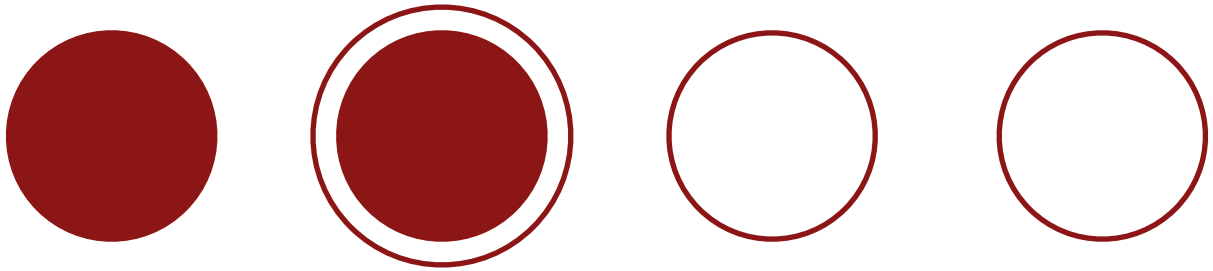

Mathematical Formulation

The Physical Problem

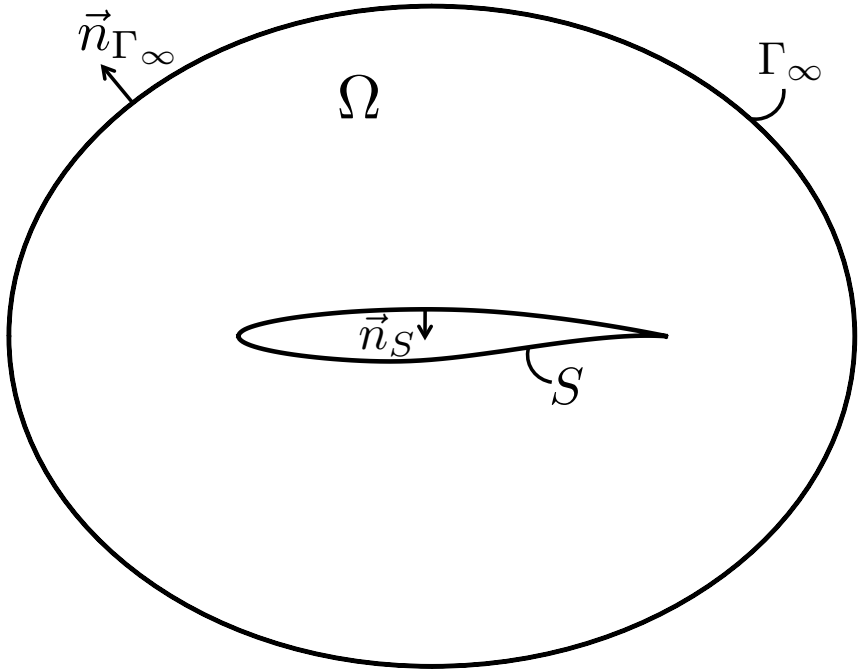

- Unsteady, Compressible Navier-Stokes Equations
- Arbitrary Lagrangian-Eulerian (ALE) formulation
- Arbitrary source term

$$
\begin{cases}
\mathcal{R}(U) = \frac{\partial U}{\partial t} + \nabla \cdot \vec{F}^c_{ale} - \nabla \cdot (\mu^1_{tot} \vec{F}^{v1} + \mu^2_{tot} \vec{F}^{v2}) - \mathcal{Q} = 0, & \text{in } \Omega, \quad t > 0 \\
\vec{v} = \vec{u}_\Omega, & \text{on } S, \\
\partial_n T = 0, & \text{on } S, \\
(W)_+ = W_\infty, & \text{on } \Gamma_\infty,
\end{cases}
$$

where…

$$U = \left\{ \begin{array}{c} \rho \\ \rho\vec{v} \\ \rho E \end{array} \right\}, \vec{F}^c_{ale} = \left\{ \begin{array}{c} \rho(\vec{v} - \vec{u}_\Omega) \\ \rho\vec{v} \otimes (\vec{v} - \vec{u}_\Omega) + \bar{\bar{I}}p \\ \rho E(\vec{v} - \vec{u}_\Omega) + p\vec{v} \end{array} \right\},$$

$$\vec{F}^{v1} = \left\{ \begin{array}{c} \cdot \\ \bar{\bar{\tau}} \\ \bar{\bar{\tau}} \cdot \vec{v} \end{array} \right\}, \vec{F}^{v2} = \left\{ \begin{array}{c} \cdot \\ \cdot \\ c_p \nabla T \end{array} \right\}, \mathcal{Q} = \left\{ \begin{array}{c} q_\rho \\ \vec{q}_{\rho\vec{v}} \\ q_{\rho E} \end{array} \right\},$$

$$\bar{\bar{\tau}} = \nabla\vec{v} + \nabla\vec{v}^{\mathsf{T}} - \frac{2}{3}\bar{\bar{I}}(\nabla \cdot \vec{v}), \; p = (\gamma - 1)\rho\left[E - \frac{1}{2}(\vec{v} \cdot \vec{v})\right], \; T = \frac{p}{\rho R}.$$

With a one-equation turbulence model for the eddy viscosity, the unsteady Reynolds-averaged Navier-Stokes (URANS) equations recovered with

$$\mu^1_{tot} = \mu_{dyn} + \mu_{tur}, \qquad \mu^2_{tot} = \frac{\mu_{dyn}}{Pr_d} + \frac{\mu_{tur}}{Pr_t}$$

The Optimization Problem

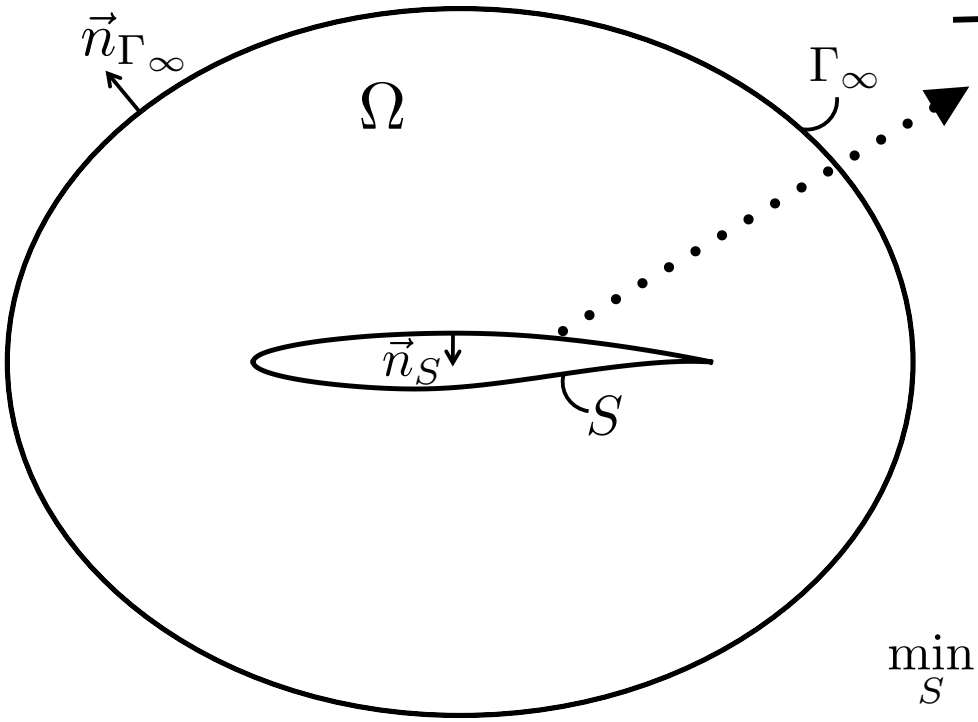

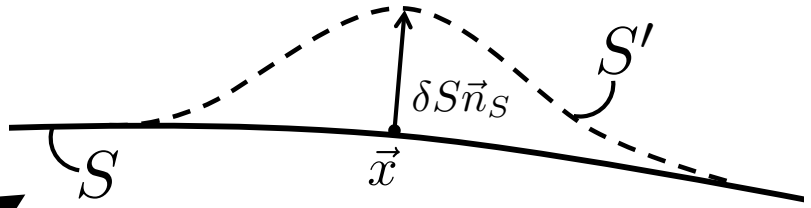

$$S' = \{\vec{x} + \delta S(\vec{x})\, \vec{n}(\vec{x}) : \ \vec{x} \in S\}$$

The time-dependent optimal shape design problem expressed as a PDE-constrained optimization problem:

$$\min_S J(S) = \frac{1}{\mathbb{T}} \int_{t_o}^{t_f} \int_S j(\vec{f}, T, \partial_n T, \vec{n})\, ds\, dt$$

$$\text{subject to: } \mathcal{R}(U) = 0$$

For efficient gradient-based optimization, let's apply the **adjoint approach** to this problem.

$$\mathbb{T} = t_f - t_o$$

Unsteady Continuous Adjoint

1. Form the Lagrangian

$$\mathcal{J} = \frac{1}{\mathbb{T}} \int_{t_o}^{t_f} \int_S j(\vec{f}, T, \partial_n T, \vec{n}) \, ds \, dt - \frac{1}{\mathbb{T}} \int_{t_o}^{t_f} \int_\Omega \Psi^\mathsf{T} \mathcal{R}(U) \, d\Omega \, dt$$

2. Take first variation w.r.t. infinitesimal surface perturbations with simplifications from differential geometry

$$\delta\mathcal{J} = \delta J - \frac{1}{\mathbb{T}} \int_{t_o}^{t_f} \int_\Omega \Psi^\mathsf{T} \delta\mathcal{R}(U) \, d\Omega \, dt$$

3. Introduce linearized N-S system (with frozen viscosity assumption)

$$\begin{cases} \delta\mathcal{R}(U) = \frac{\partial}{\partial t}(\delta U) + \nabla \cdot \left(\vec{A}^c - \bar{\bar{I}} \vec{u}_\Omega - \mu_{tot}^k \vec{A}^{vk} \right) \delta U - \nabla \cdot \mu_{tot}^k \bar{\bar{D}}^{vk} \delta(\nabla U) - \frac{\partial \mathcal{Q}}{\partial U} \delta U = 0 & \text{in } \Omega, \ t > 0 \\ \delta\vec{v} = -\partial_n(\vec{v} - \vec{u}_\Omega)\delta S & \text{on } S, \\ \partial_n(\delta T) = \nabla T \cdot \nabla_S(\delta S) - \partial_n^2(T)\delta S & \text{on } S, \\ (\delta W)_+ = 0 & \text{on } \Gamma_\infty \end{cases}$$

4. Integrate by parts (once each for time derivative term & convective term, twice for viscous terms)

$$\delta\mathcal{J} = \delta J - \frac{1}{\mathbb{T}}\int_{\Omega}\left[\Psi^{\mathsf{T}}\delta U\right]_{t_o}^{t_f} d\Omega - \frac{1}{\mathbb{T}}\int_{t_o}^{t_f}\int_{\partial\Omega}\left(B_1 - B_2 + B_3\right) ds\, dt$$
$$- \frac{1}{\mathbb{T}}\int_{t_o}^{t_f}\int_{\Omega}\left[-\frac{\partial\Psi^{\mathsf{T}}}{\partial t} - \nabla\Psi^{\mathsf{T}}\cdot\left(\vec{A}^{c} - \bar{\bar{I}}\vec{u}_{\Omega} - \mu_{tot}^{k}\vec{A}^{vk}\right) - \nabla\cdot\left(\nabla\Psi^{\mathsf{T}}\cdot\mu_{tot}^{k}\bar{\bar{D}}^{vk}\right) - \Psi^{\mathsf{T}}\frac{\partial\mathcal{Q}}{\partial U}\right]\delta U\, d\Omega\, dt$$

5. Evaluate the boundary integral terms by hand given our knowledge of the equations and linearized boundary conditions

$$B_1 = \Psi^{\mathsf{T}}\left(\vec{A}^{c} - \bar{\bar{I}}\vec{u}_{\Omega}\right)\delta U\cdot\vec{n}$$
$$B_2 = \Psi^{\mathsf{T}}\mu_{tot}^{k}\vec{A}^{vk}\delta U\cdot\vec{n} + \Psi^{\mathsf{T}}\mu_{tot}^{k}\bar{\bar{D}}^{vk}\cdot\nabla(\delta U)\cdot\vec{n}$$
$$B_3 = \nabla\Psi^{\mathsf{T}}\cdot\mu_{tot}^{k}\bar{\bar{D}}^{vk}\delta U\cdot\vec{n}$$

Unsteady Continuous Adjoint

6. Domain integrals form the adjoint equations along with admissible BCs identified from previous step - satisfying these removes dependence on flow variations

$$\begin{cases} -\frac{\partial \Psi^\mathsf{T}}{\partial t} - \nabla\Psi^\mathsf{T} \cdot \left(\vec{A}^c - \bar{\bar{I}}\vec{u}_\Omega - \mu^k_{tot}\vec{A}^{vk}\right) - \nabla \cdot \left(\nabla\Psi^\mathsf{T} \cdot \mu^k_{tot}\bar{\bar{D}}^{vk}\right) - \Psi^\mathsf{T}\frac{\partial \mathcal{Q}}{\partial U} = 0, & \text{in } \Omega, \quad t > 0 \\ \vec{\varphi} = \vec{d} - \psi_{\rho E}\,\vec{v}, & \text{on } S, \\ \partial_n(\psi_{\rho E}) = 0, & \text{on } S, \\ \Psi = 0, & \text{in } \Omega, \quad t = t_o, t_f. \end{cases}$$

7. After satisfying the adjoint equations, recover the expression for the surface sensitivity from remaining boundary terms

$$\begin{aligned} \delta\mathcal{J} = \frac{1}{\mathbb{T}}\int_{t_o}^{t_f}\int_S \{\vec{d}\cdot[\vec{q}_{\rho\vec{v}} - \partial_t(\rho\vec{v})] + \nabla\vec{d} : (\bar{\bar{I}}p - \bar{\bar{\sigma}}) - (\bar{\bar{I}}p - \bar{\bar{\sigma}})\cdot\vec{n}\cdot\partial_n\vec{d} + \vartheta\partial_n(\vec{v} - \vec{u}_\Omega)\cdot\vec{n} \\ - \psi_{\rho E}\partial_n(\vec{v} - \vec{u}_\Omega)\cdot\bar{\bar{\sigma}}\cdot\vec{n} + \vec{n}\cdot\left(\bar{\bar{\Sigma}}^{\varphi} + \bar{\bar{\Sigma}}^{\psi_{\rho E}}\right)\cdot\partial_n(\vec{v} - \vec{u}_\Omega) - \mu^2_{tot}c_p\nabla_S(\psi_{\rho E})\cdot\nabla_S(T) \\ - \psi_{\rho E}\left[p(\nabla\cdot\vec{v}) - \bar{\bar{\sigma}} : \nabla\vec{v} + \partial_t(\rho E) + (\vec{q}_{\rho\vec{v}} - \partial_t(\rho\vec{v}))\cdot\vec{v} - q_{\rho E}\right]\}\delta S\, ds\, dt, \\ = \frac{1}{\mathbb{T}}\int_{t_o}^{t_f}\int_S \left\{\frac{\partial\mathcal{J}}{\partial S}\right\}\delta S\, ds\, dt \end{aligned}$$

Surface Sensitivity

Instantaneous Flow Solution

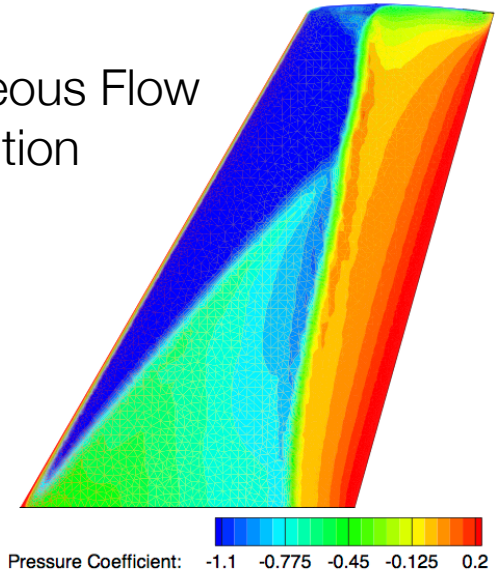

Instantaneous Sensitivity for a Drag Objective

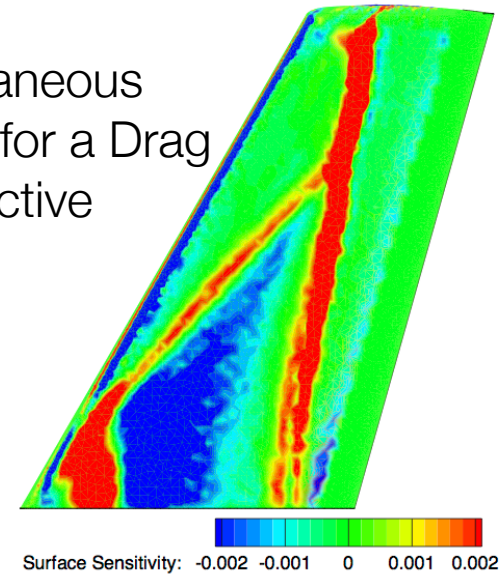

- Key result of the continuous adjoint derivation
- Measures change in the objective function w.r.t. small perturbations in the local normal direction
- Computed at each surface mesh node (at each time step!) at negligible cost from the flow and adjoint variables on the surface
- Gradient expression is a surface integral: **no dependence on the volume mesh**
- Offers physical insight and designer intuition

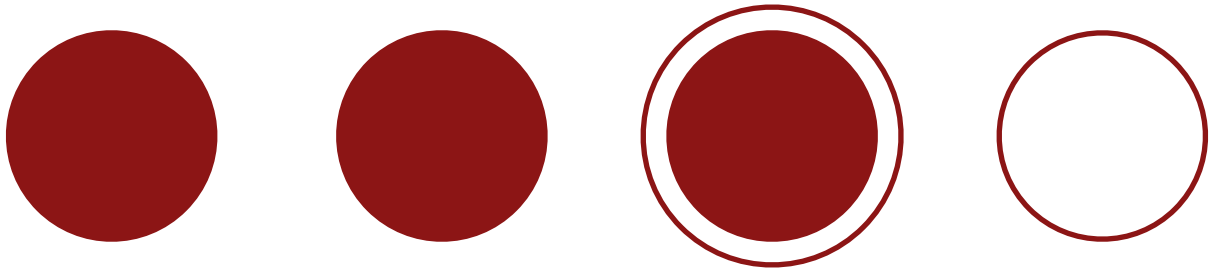

Numerical Results

Optimal Shape Design

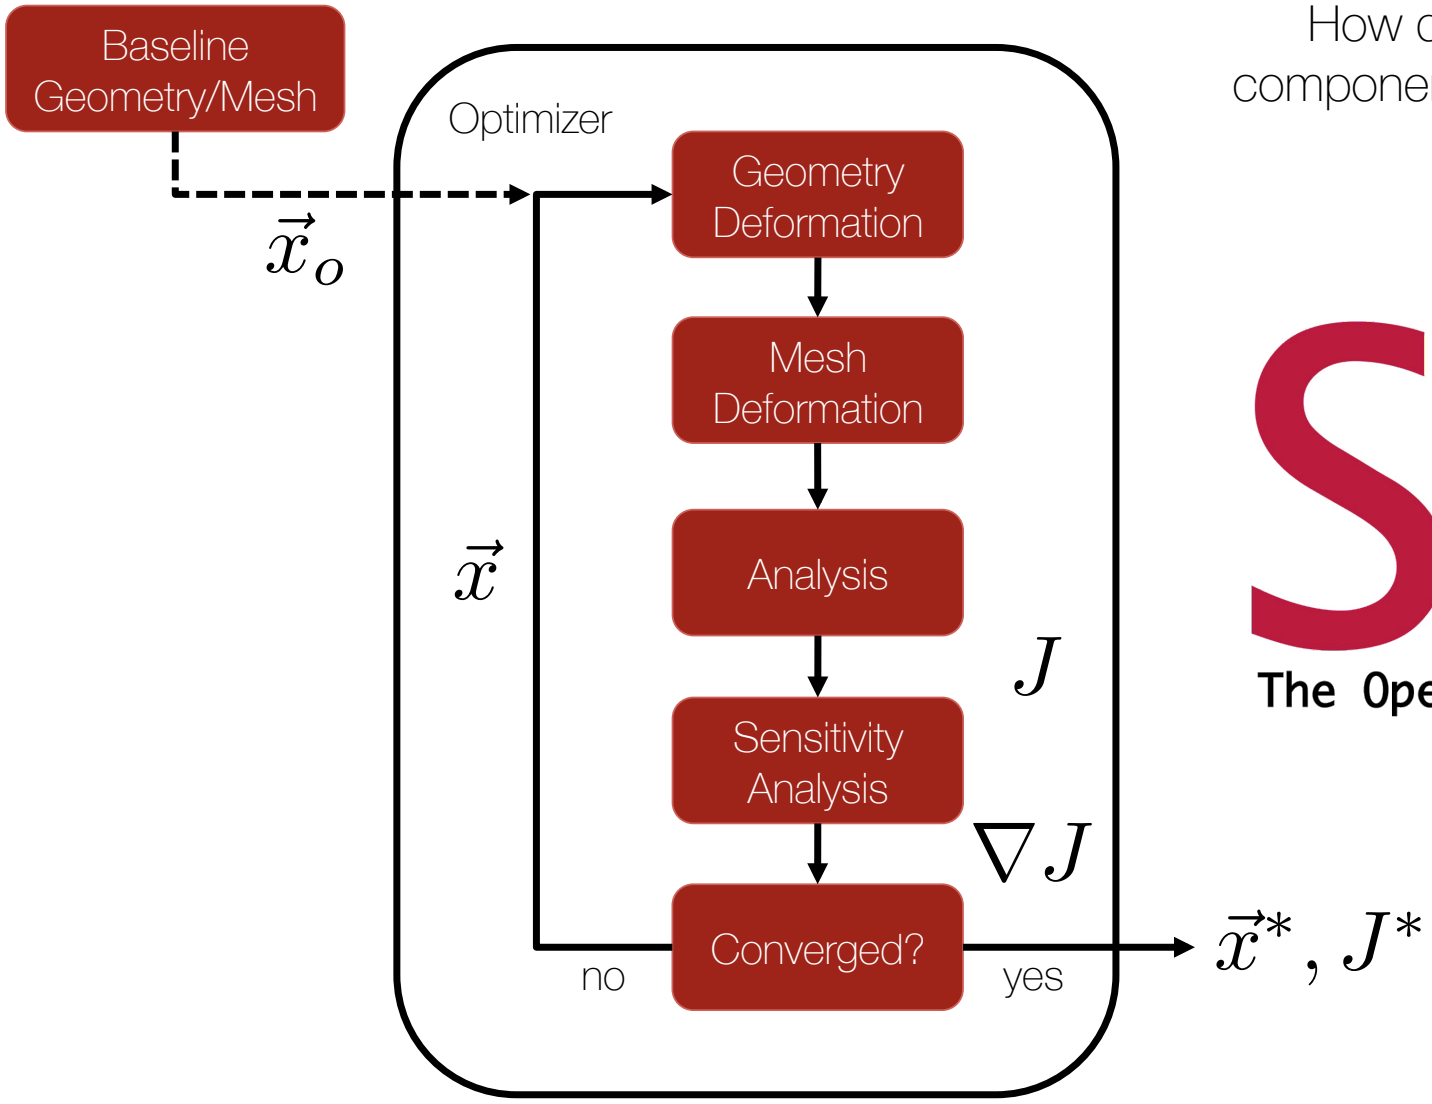

How do we implement the components of the optimal shape design loop?

SU2

The Open-Source CFD Code

Geometry Parameterization

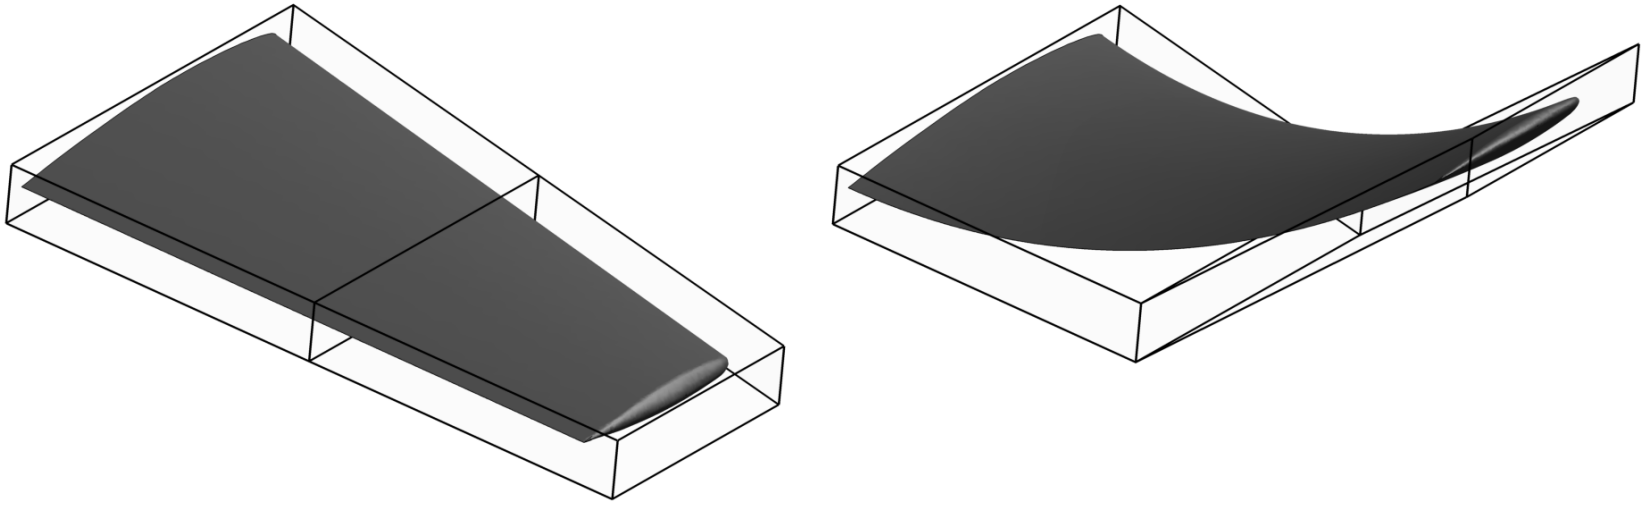

- In 3-D, parameterize by a **Free-Form Deformation** (FFD) approach
- Origins in the computer graphics industry
- Encapsulate geometry in a bounding box and create a mapping between the FFD control points and the mesh surface nodes (parameterize as a Bézier solid)
- FFD control points become the DVs with the surface inheriting a smooth deformation

Mesh Deformation

Solve the **linear elasticity equations** on the volume mesh to compute nodal displacements.

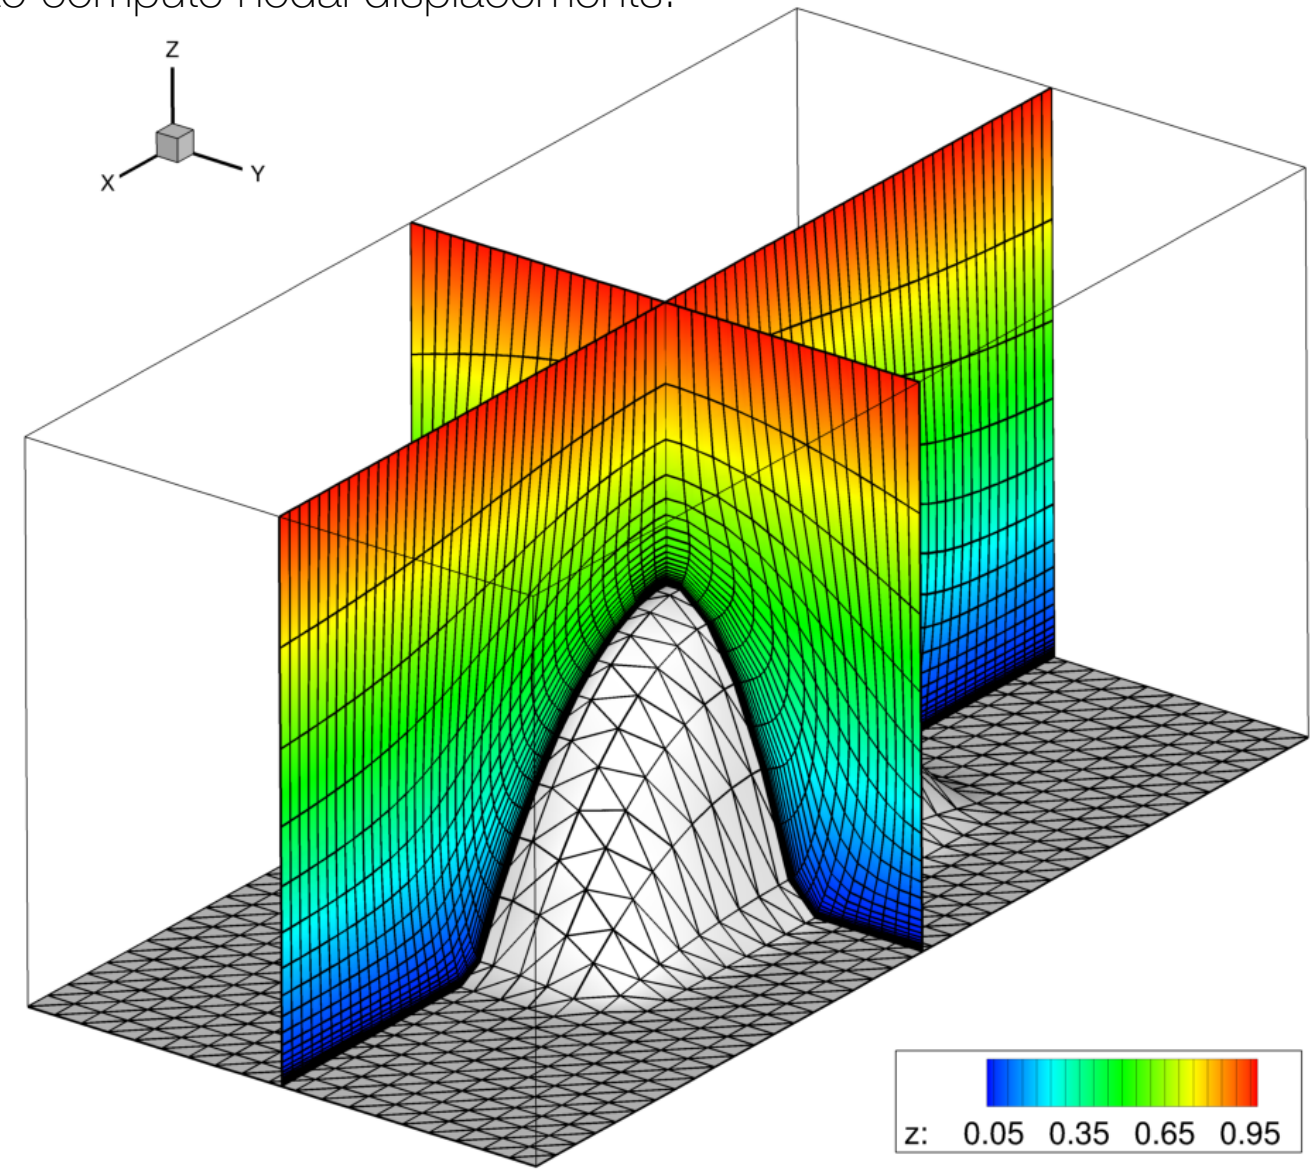

Flow/Adjoint PDE Analysis

- Unstructured meshes with median-dual control volumes (vertex-based)
- Finite Volume Method with second-order schemes in space:
  - JST, Roe (+ limiting), HLLC, AUSM for convective fluxes
  - Average of gradients for viscous fluxes
  - Piecewise constant source terms
- Similar schemes for spatially integrating the flow and adjoint, although the adjoint is treated with a non-conservative scheme
- Dual-time stepping approach for second-order accurate time integration
- Reverse time integration for the adjoint problem by writing/reading flow solutions to disk (checkpointing is possible)

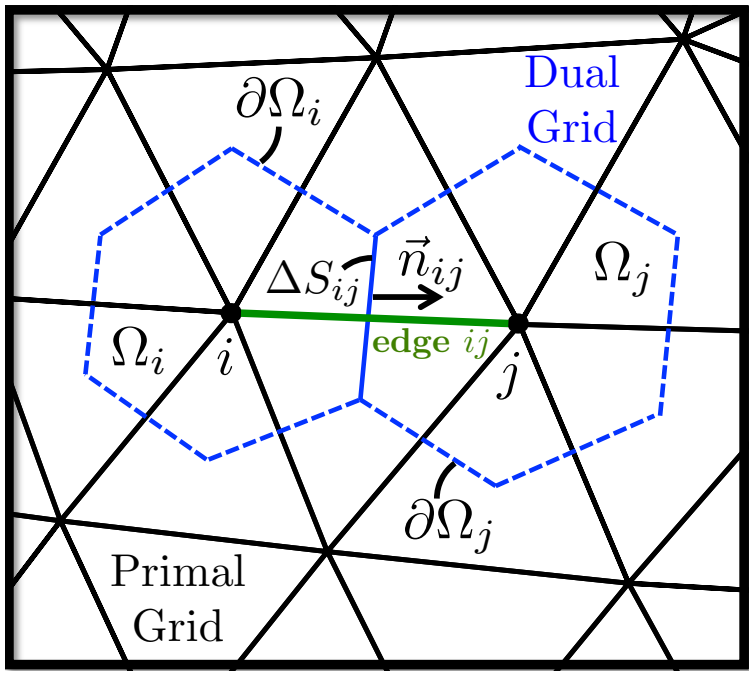

Dynamic Meshes

- Dynamic meshes are required with the ALE formulation
- Typically, either **rigidly transforming** or **dynamically deforming** grids are used
- Must also satisfy the Geometric Conservation Law (GCL) when simulating unsteady flows on dynamic grids
- The pitching wing results that follow are based on rigid mesh transformations (i.e., there is no relative motion between individual grid nodes)
- Specify the angle of attack in time and pitch the mesh as a solid body about the chosen pitching axis:

$$\alpha(t) = \alpha_o + \alpha_m \sin(\omega t)$$

Optimal Shape Design

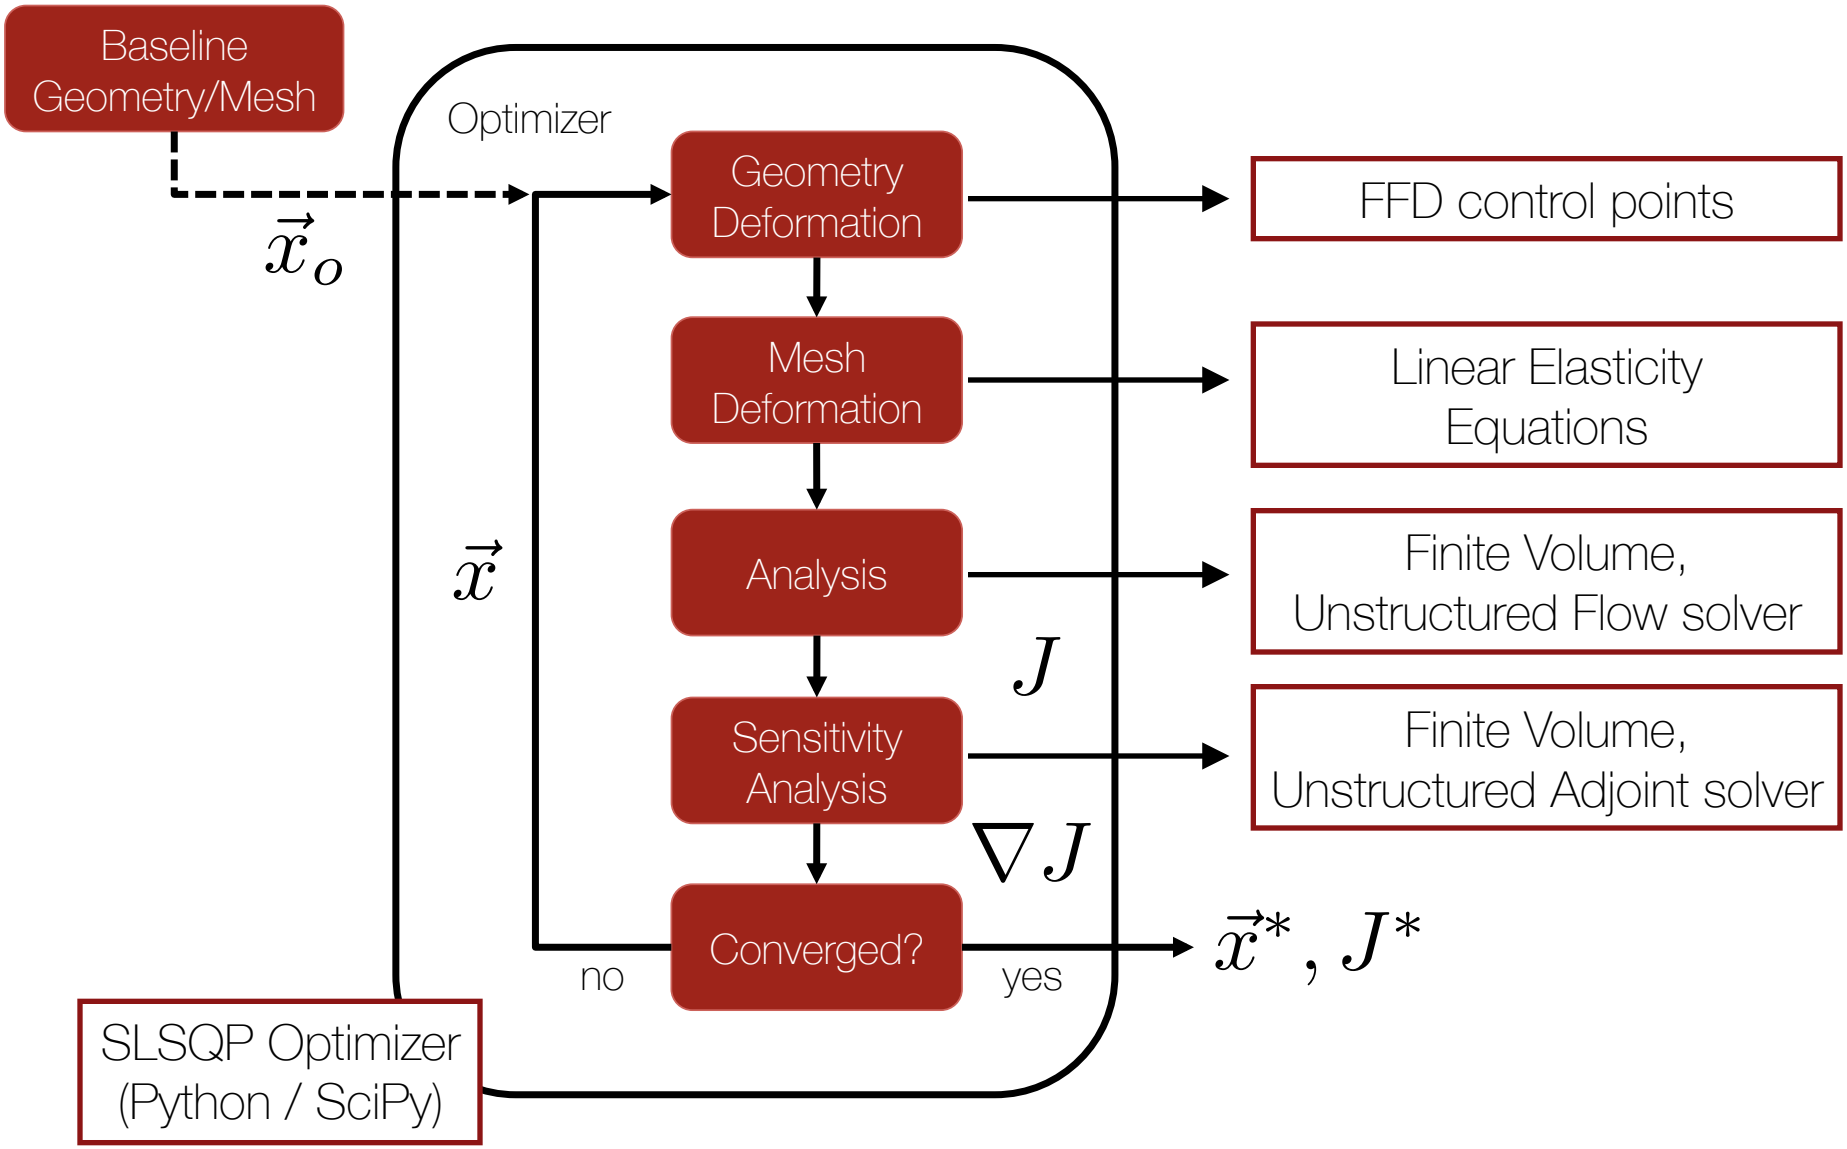

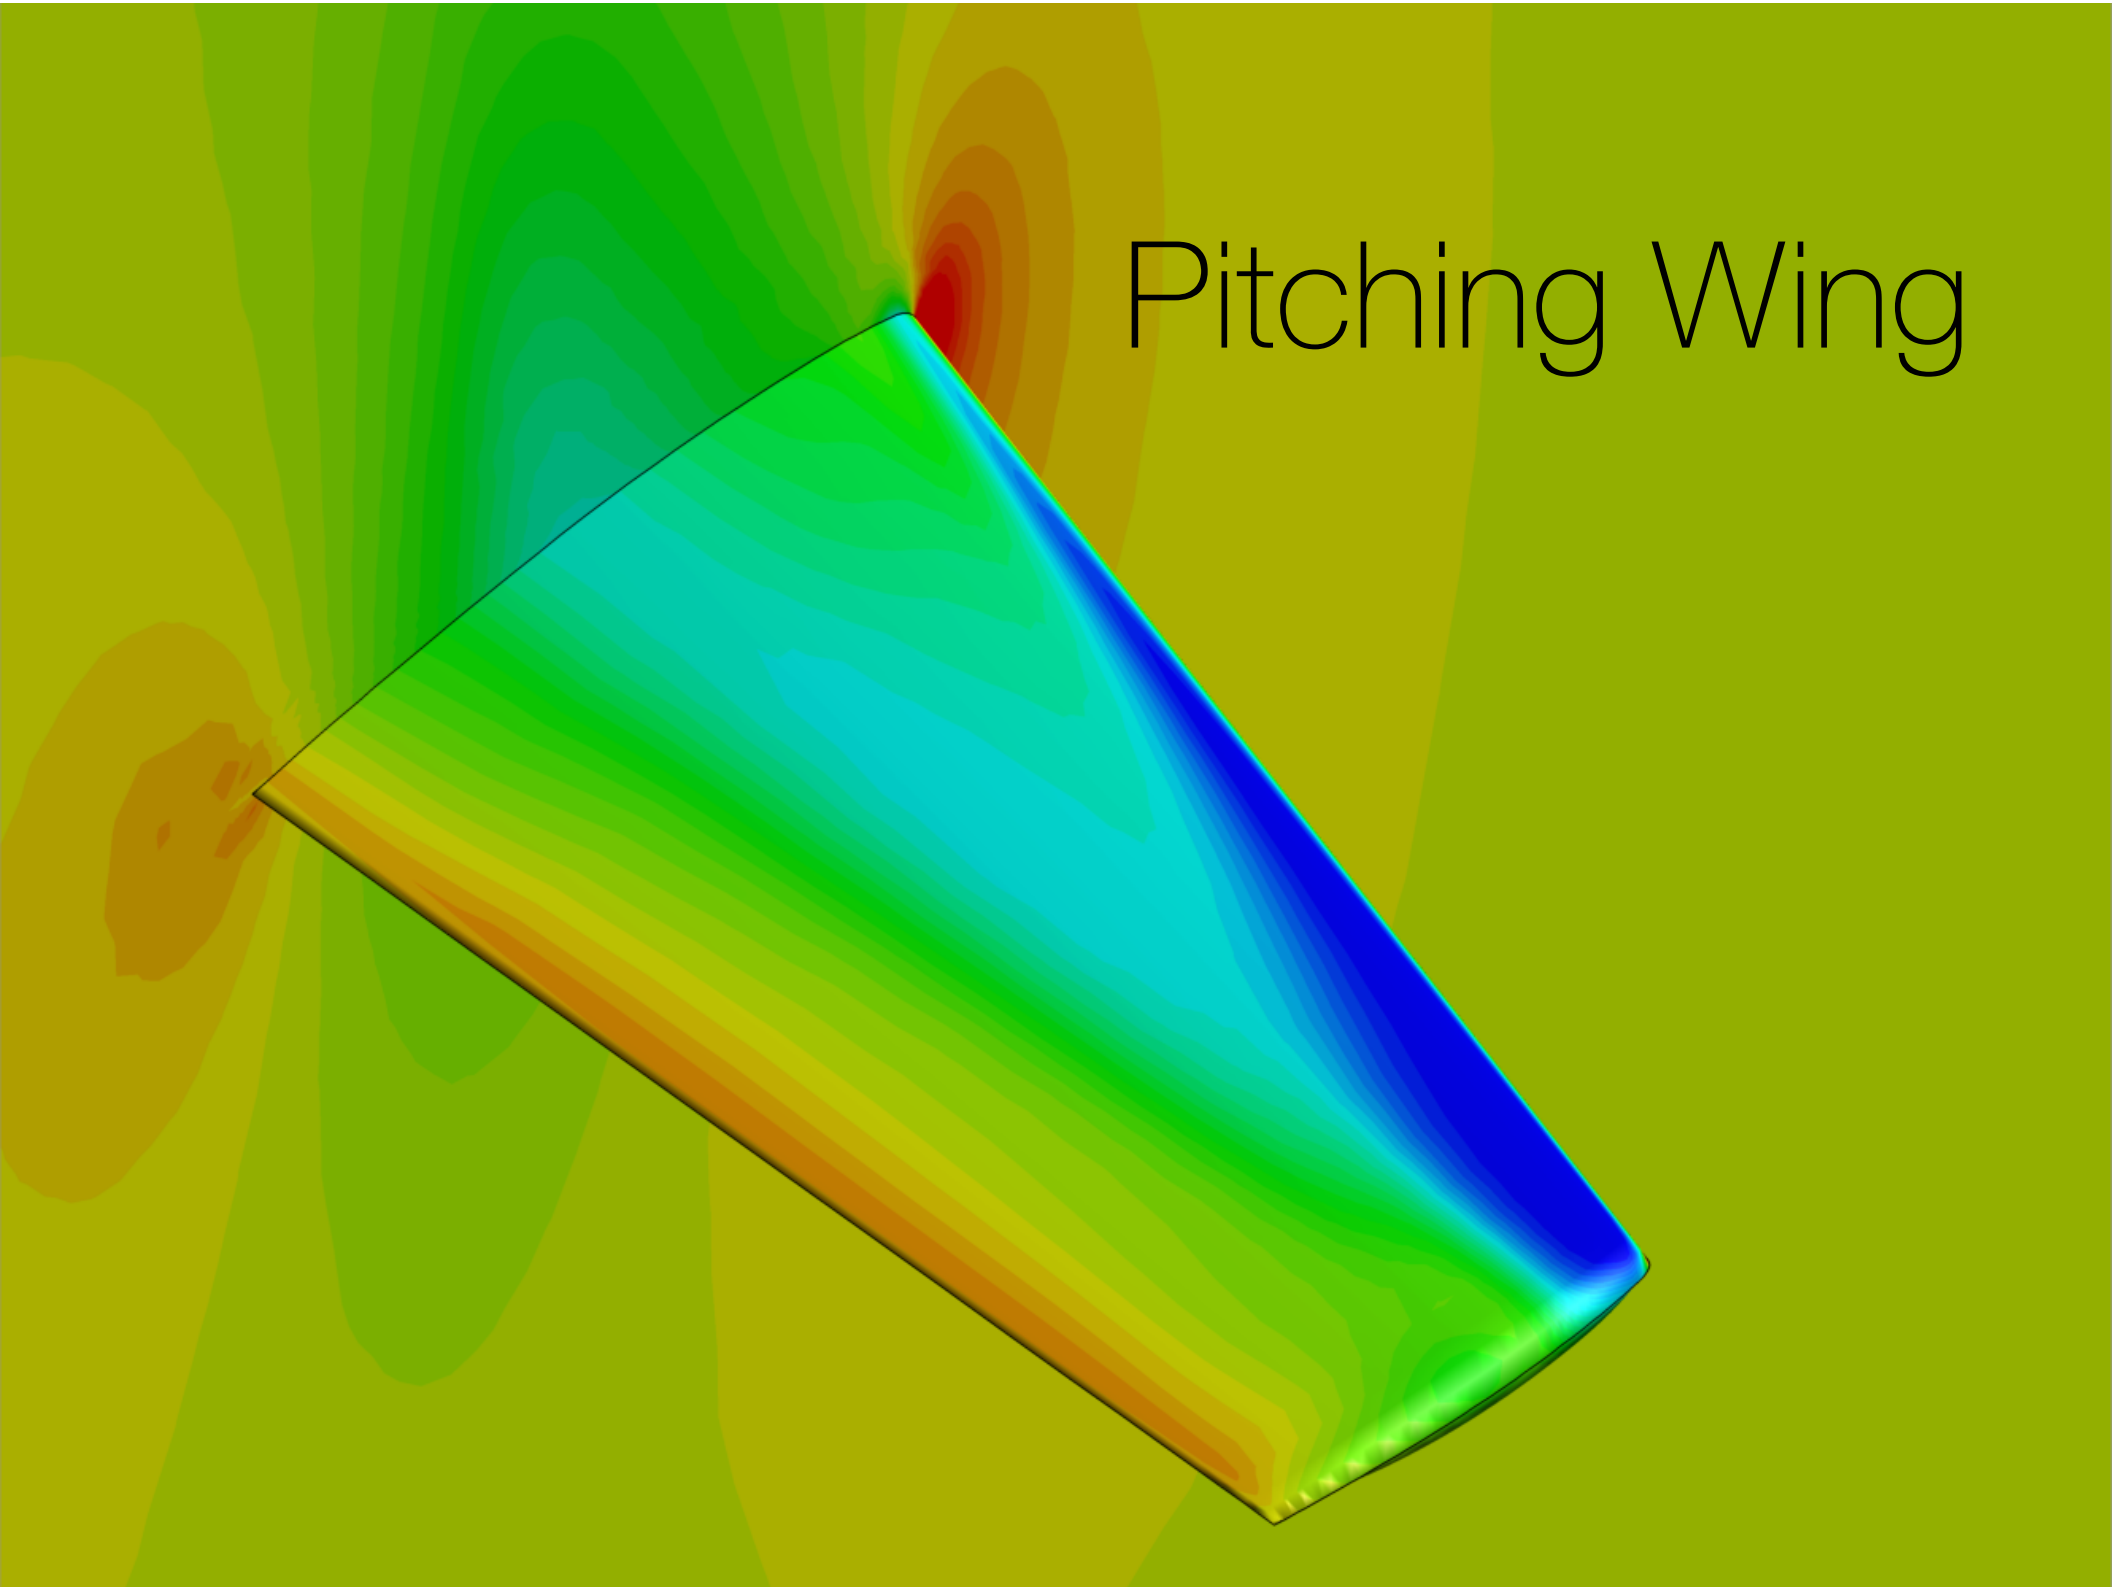
Pitching Wing

Pitching Wing

- ONERA M6 wing as the baseline geometry
- Pitching wing in transonic flow:
  - Mach = 0.8395
  - reduced frequency of 0.1682
  - mean alpha of 3.06 degrees
  - pitching amplitude of 2.5 degrees
  - Reynolds number = 11.72 million
- Pitching about the y-axis through the root quarter-chord
- 25 time steps per period for 7 periods
- RANS (S-A) equations on rigidly transforming meshes

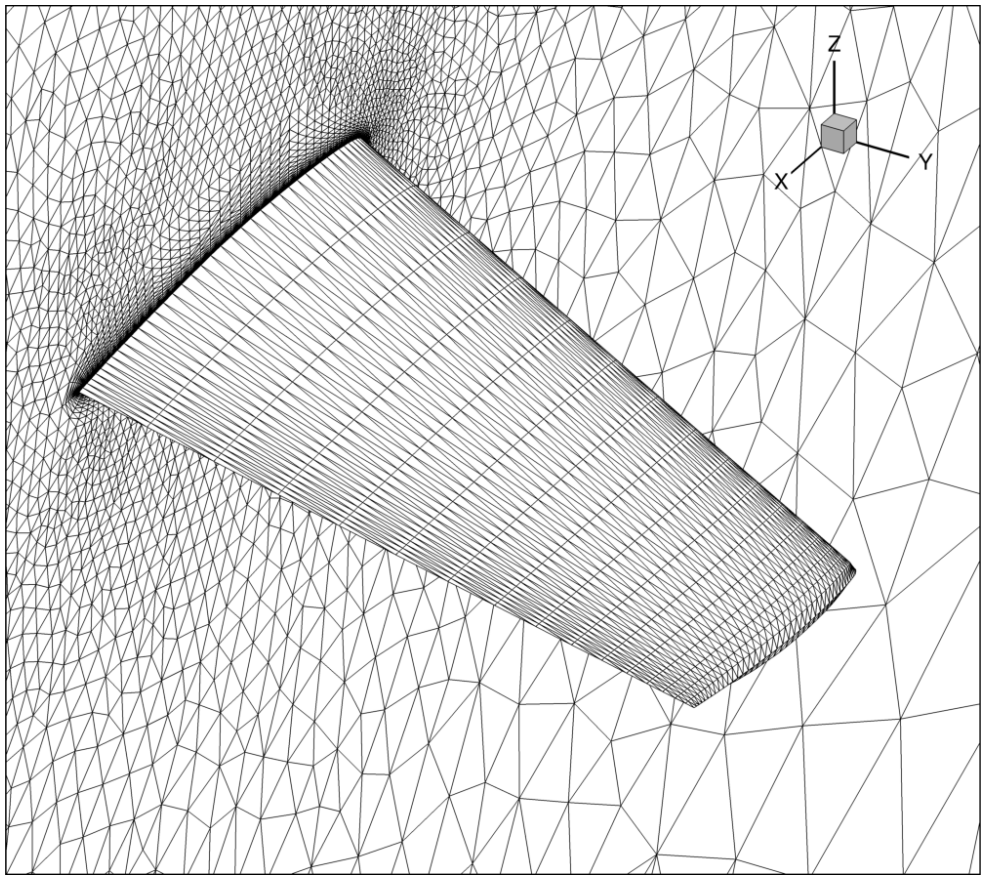

Pitching Wing

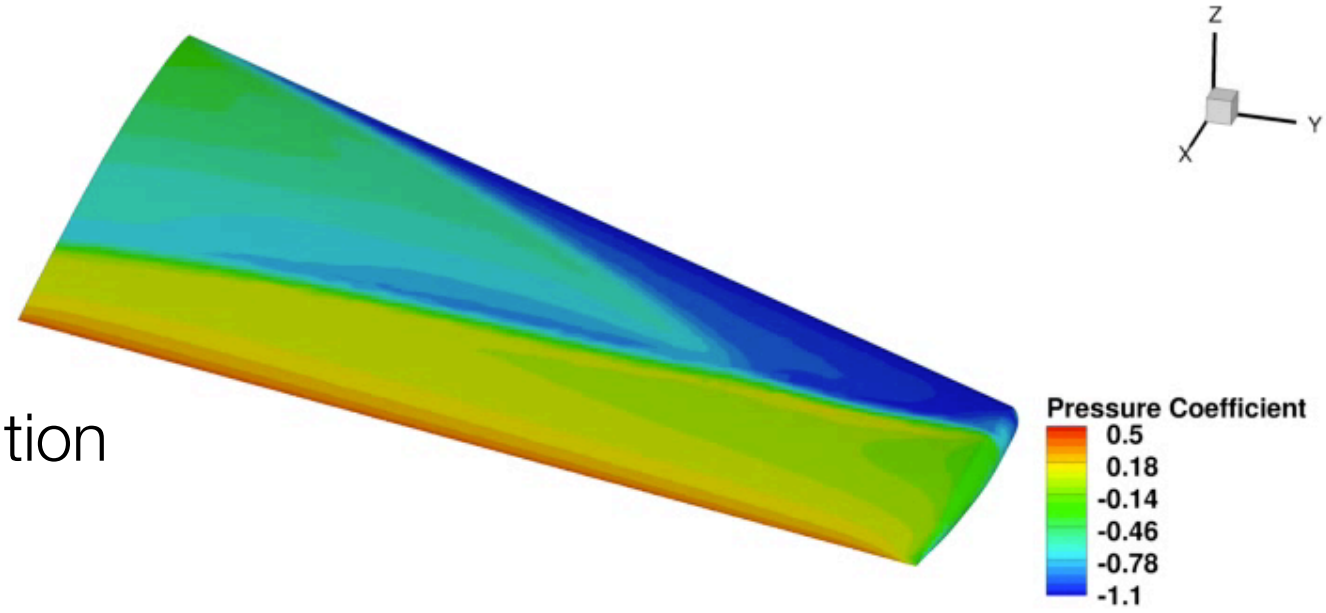

Flow Solution

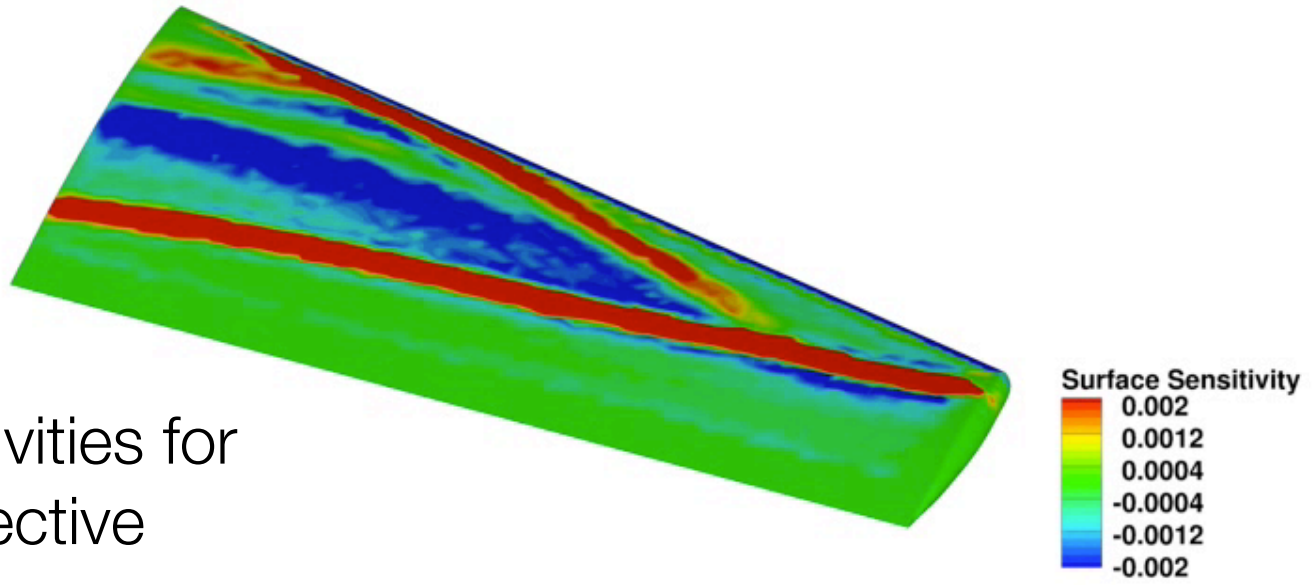

Surface Sensitivities for a Drag Objective

(Euler Results for Illustration)

Animation

Pitching Wing

Constrained, Time-averaged Drag Minimization with Lift and Maximum Thickness Constraints

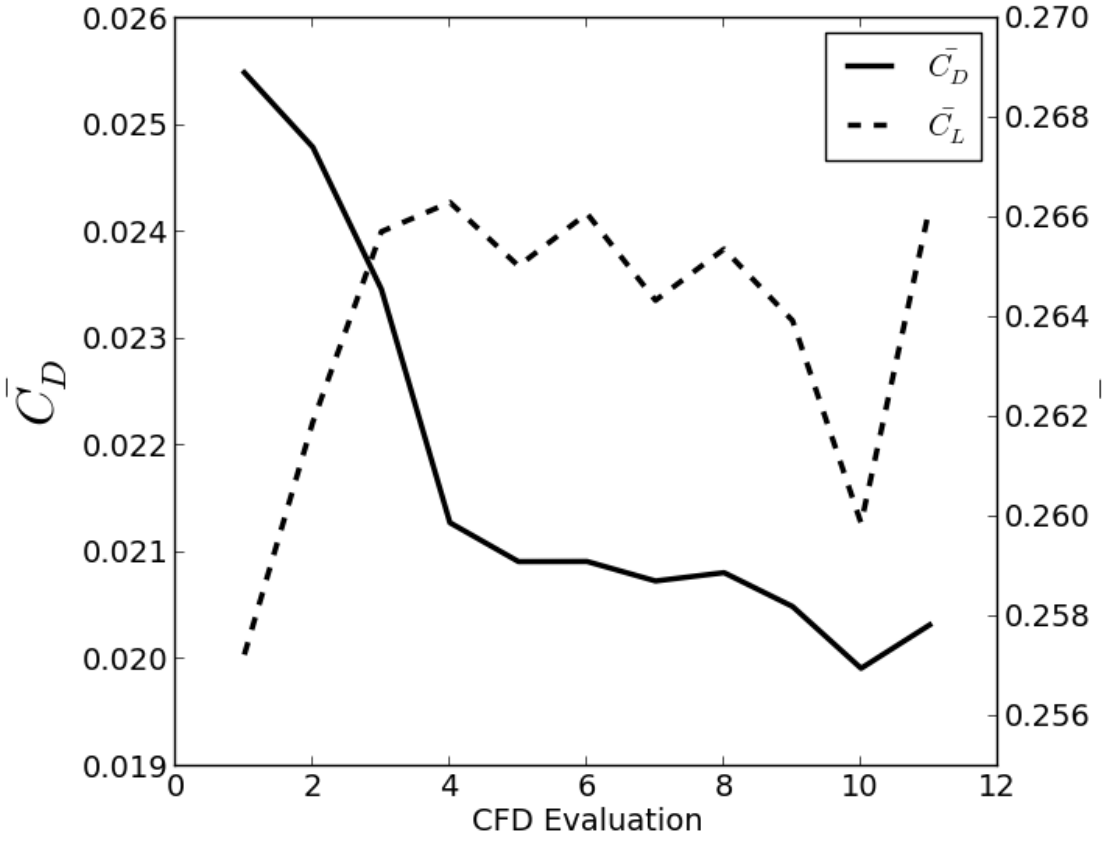

20.3 % Reduction in Time-averaged Drag

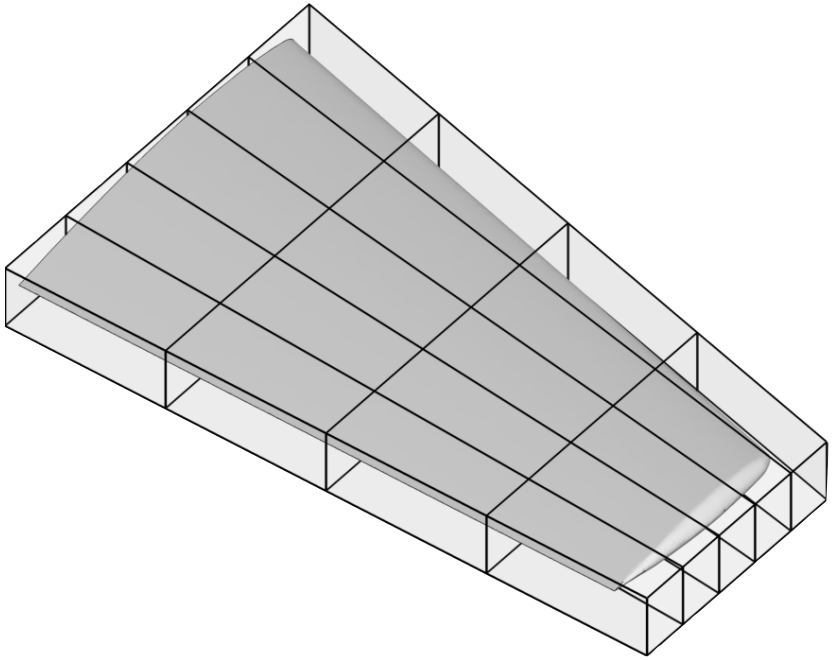

FFD Control Point Variables.
Movement allowed in +/- z-direction.

Drag Coefficient Time Histories

Large drag reductions near the incidence of max drag

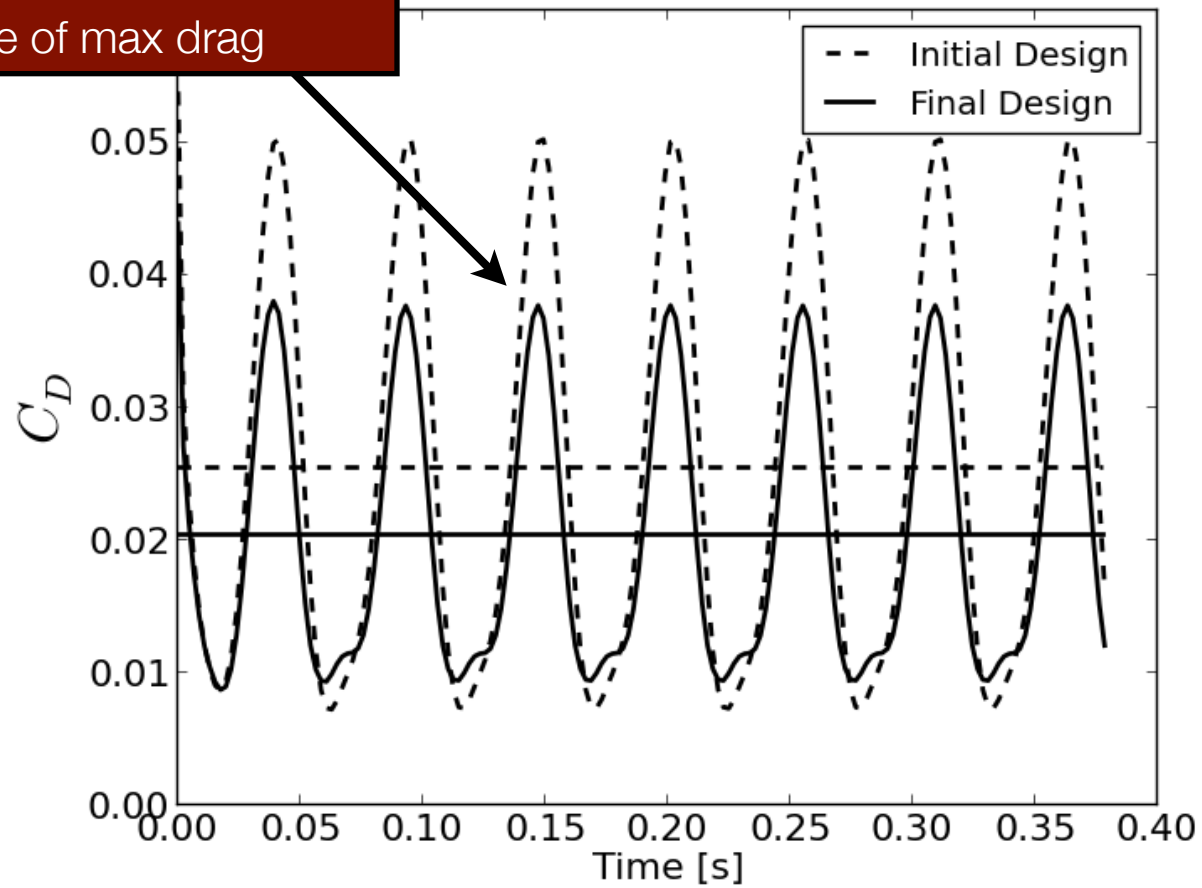

Pitching Wing

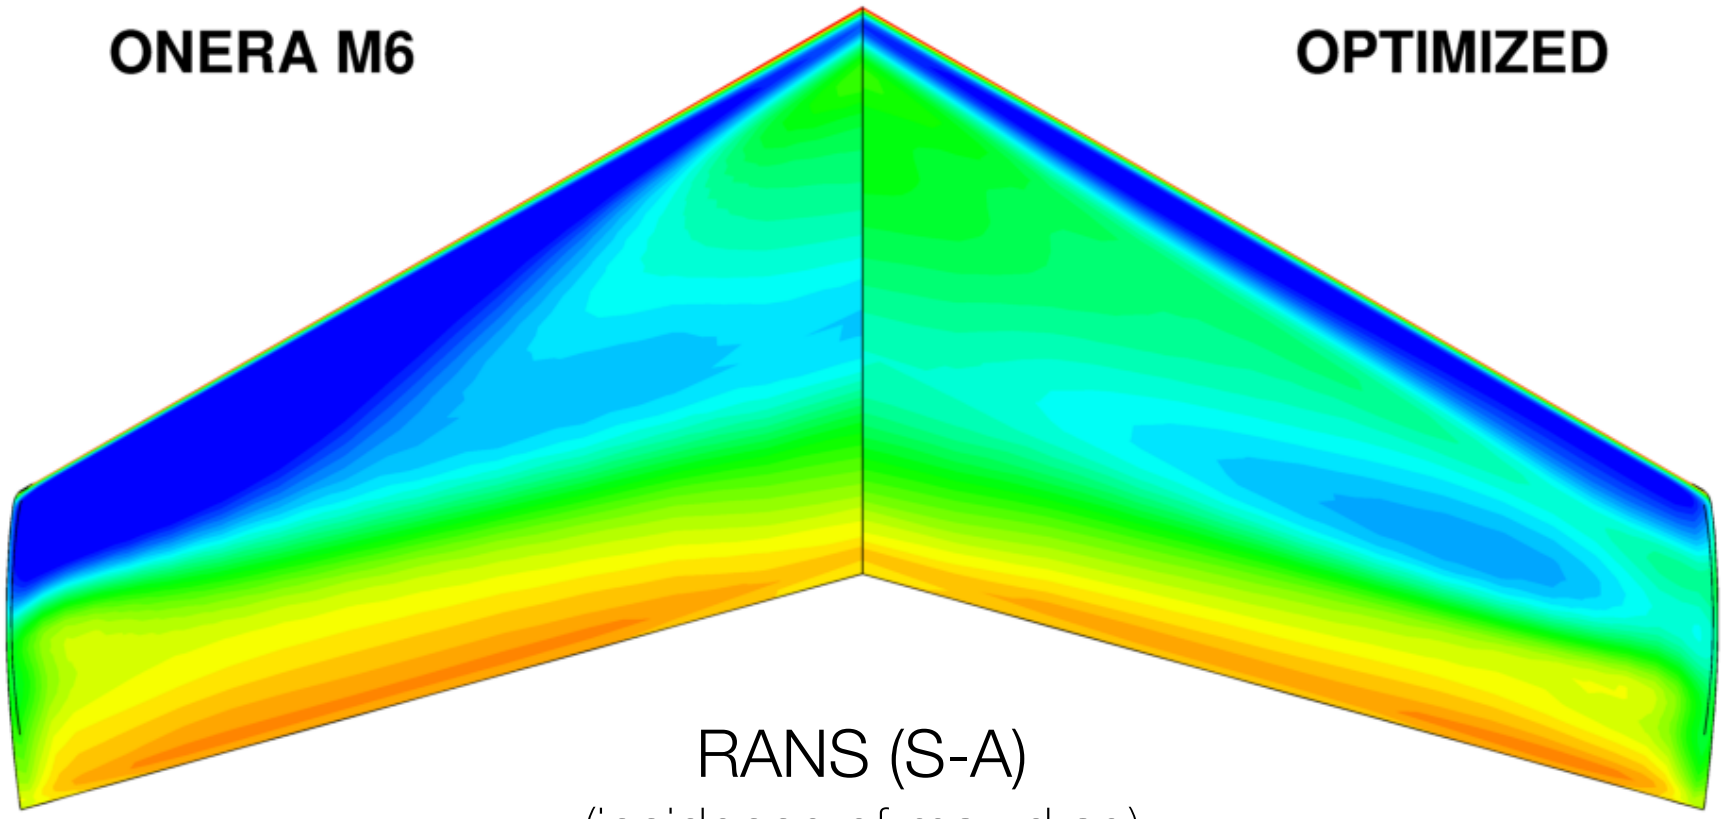

RANS (S-A)

(incidence of max drag)

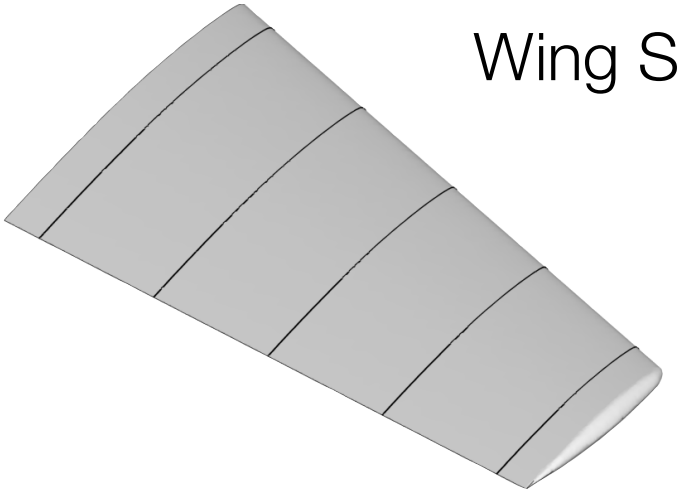

Wing Section Shape Comparison

RANS

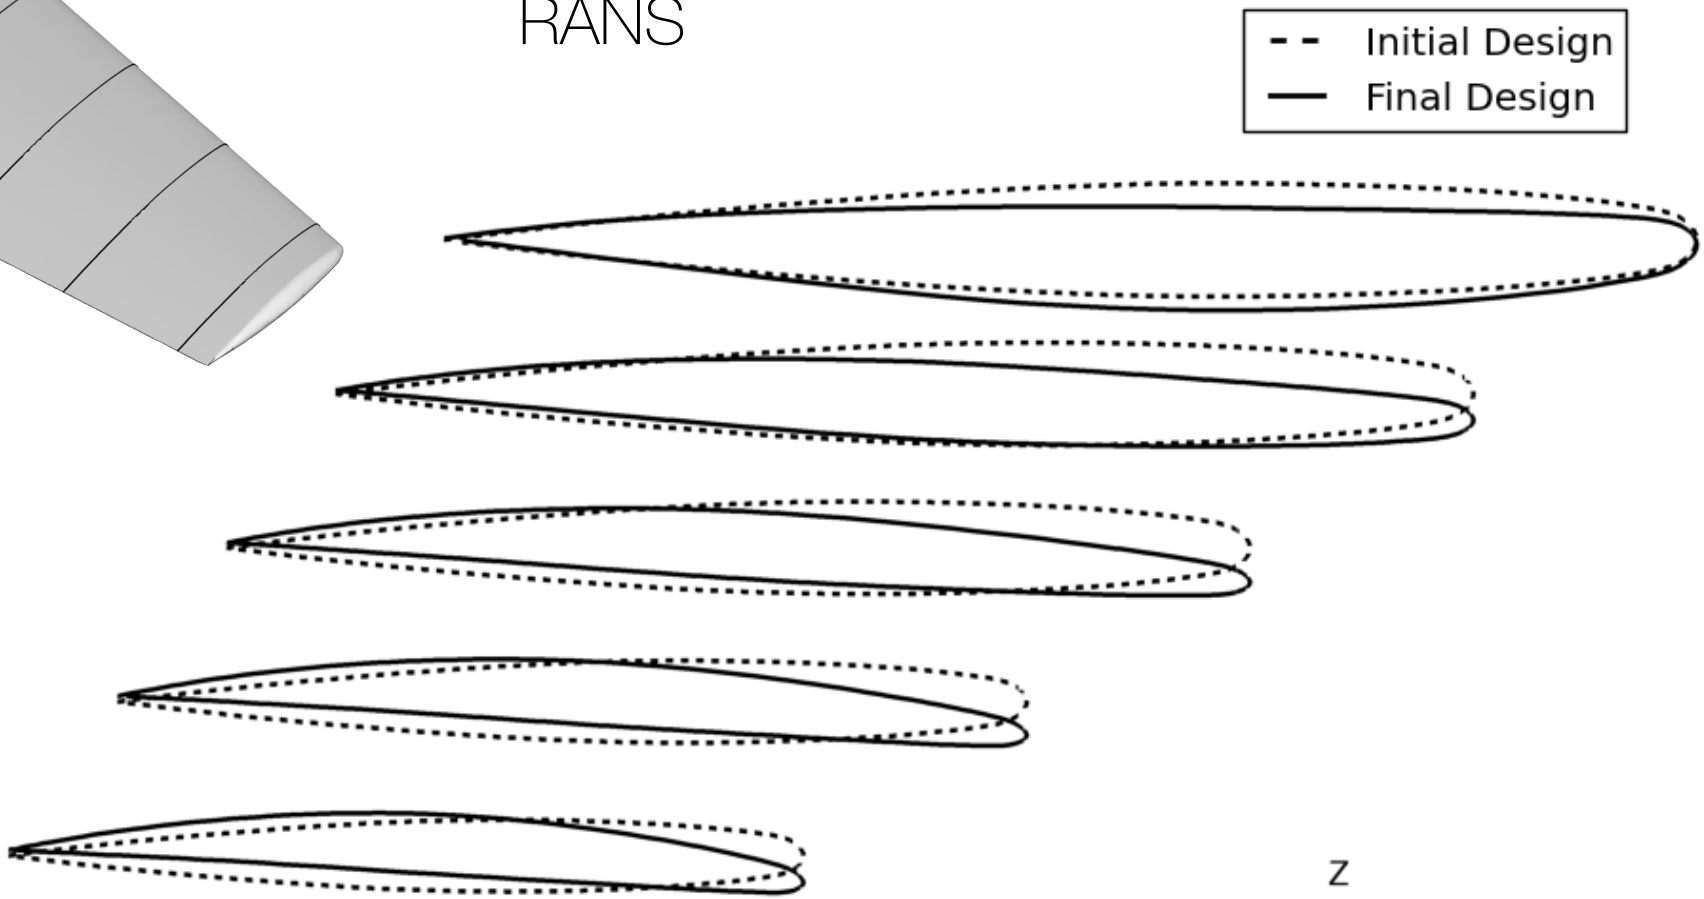

Unsteady Adjoitn Conclusions

- Developed the first time-accurate continuous adjoint for the unsteady, compressible Navier-Stokes equations (ALE form) using shape calculus
  - Particular emphasis on boundary conditions, surface sensitivity, and simplifications (more details in article/thesis)
- Formulation contains all of the complexity needed for treating large-scale problems of industrial interest for unsteady flows on dynamic meshes (with a big enough computer…). **See next part of talk for a number of future considerations.**
- Unlocks other time-dependent optimization problems such as aeroelastics, aeroacoustics, flow control, etc.

Economon, T. D., Palacios, F., Alonso, J. J., “An Unsteady Continuous Adjoint Approach for Aerodynamic Design on Dynamic Meshes," AIAA Journal, accepted for publication, 2015.

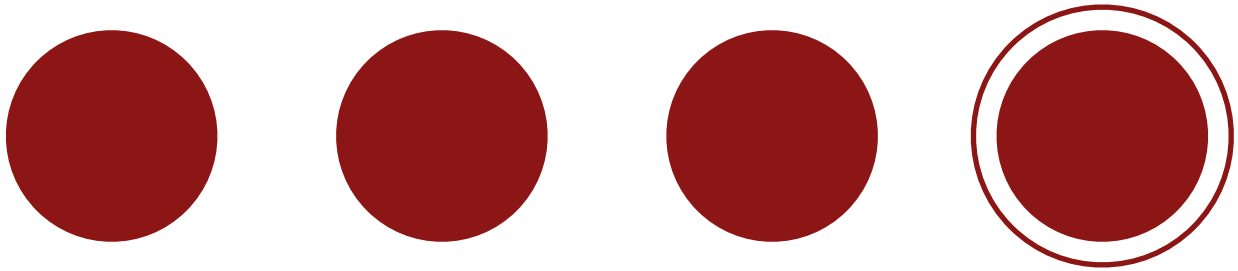

Future Directions for Adjoints

(Adjoint) State of the Union

- The solution of the RANS equations around complete aircraft configurations is an **everyday occurrence in the aerospace industry**.
- **Gradient-based optimization methods coupled with adjoint techniques** for sensitivity analysis offer the most efficient process for detailed shape design.
- However, their **adoption for realistic problems in industrial settings has lagged due to several main difficulties:** robustness, and an inability to handle large, complex meshes.
- This situation merits a renewed focus on techniques that offer the additional flexibility to overcome these roadblocks, while also **providing sufficient accuracy and computational performance**.
- Once the **right tool is available** the aircraft designers should develop a **methodology to use this tool** in the context of complex configurations.

Redesign of the NASA Common Research Model

The CRM is a low-wing, standard, tube-and-wing configuration. The on-design conditions are Mach 0.85 and a nominal lift of C_L=0.5 at Re = 40 million per reference chord.

- How does the adjoint perform for an industry-relevant problem in aeronautics?
- Here, we consider complex geometry and conditions requiring unstructured meshes with tens of millions of cells running on hundreds or thousands of processing elements. RANS flow with the Navier-Stokes adjoint.
- What is the status in terms of accuracy & robustness? Any opportunities?
- What about computational performance?

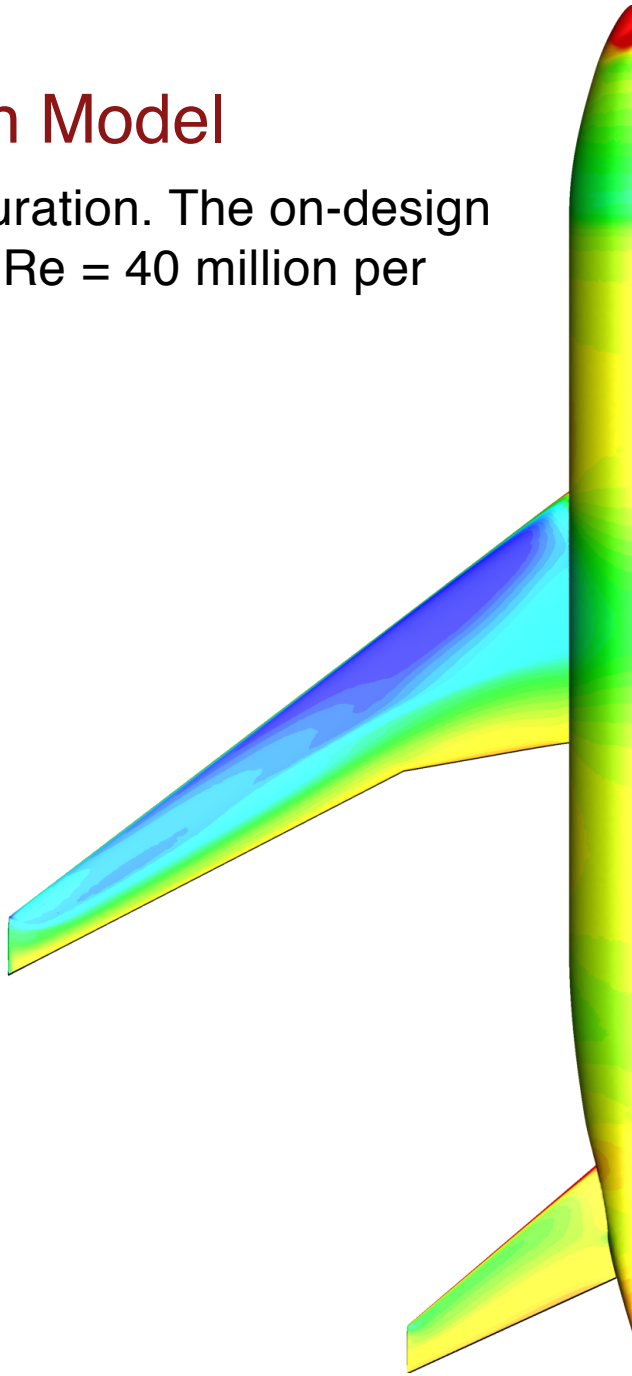

The Platform

The SU2 suite is an **open-source** collection of C++ based software for PDE analysis and PDE-constrained optimization (i.e., Computational Fluid Dynamics!).

SU2 is under active development at Stanford University in the Aerospace Design Lab (ADL) of the Department of Aeronautics and Astronautics and **now in many places around the world**.

http://su2.stanford.edu/

2012 SU² team, "Stanford University Unstructured (SU2): An open-source integrated computational environment for multi-physics simulation and design", AIAA Paper 2013-0287.
2013 SU² team, "Stanford University Unstructured (SU2): Open-source analysis and design technology for turbulent flows", AIAA Paper 2014-0243.

Surface Formulation Assessment

The expression for the total variation of the functional can be simplified as follows with $\boldsymbol{\Sigma}$ depending on the gradient of the adjoint variables.

$$\delta\mathcal{J} = \int_S (\vec{n} \cdot \bar{\Sigma}^{\varphi} \cdot \partial_n \vec{v} - \mu_{tot}^2 C_p \nabla_S \psi_5 \cdot \nabla_S T)\, \delta S\, ds$$

C. Castro, C. Lozano, F. Palacios, and E. Zuazua "A Systematic Continuous Adjoint Approach to Viscous Aerodynamic Design on Unstructured Grids", AIAA Journal, 2007.

How accurate is this formula? It is exact for Navier-Stokes problems, it does not include mesh dependent issues and it is consistent with turbulence models.

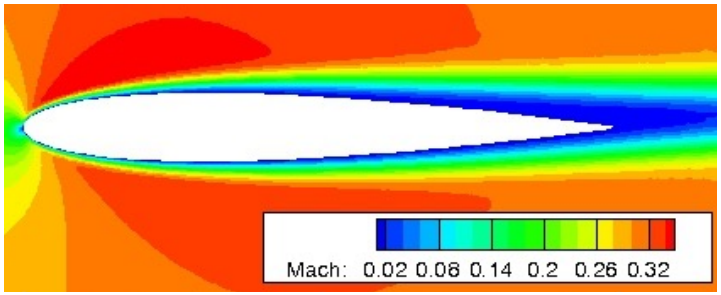

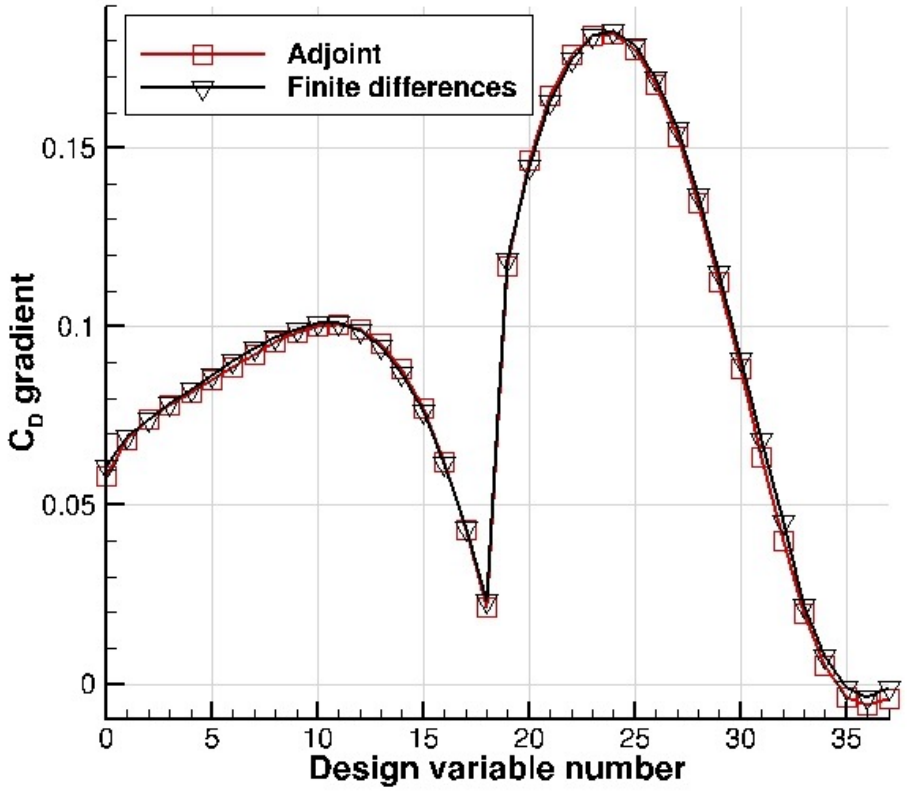

Surface Formulation Assessment (Transonic)

- We said it was exact for Navier-Stokes problems, but we assumed a smooth flow solution in the derivation… what about shocks?
- While formulations that explicitly treat shocks have been developed, in practice, it is unnecessary given that we will not see true discontinuities in the numerical solution of the primal problem.

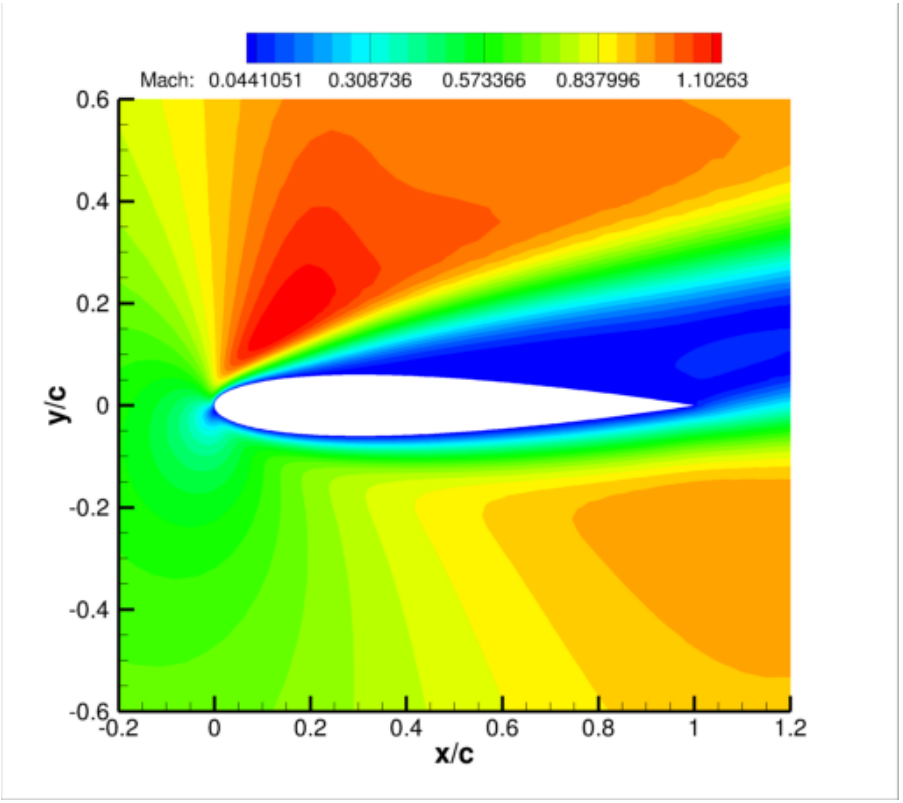

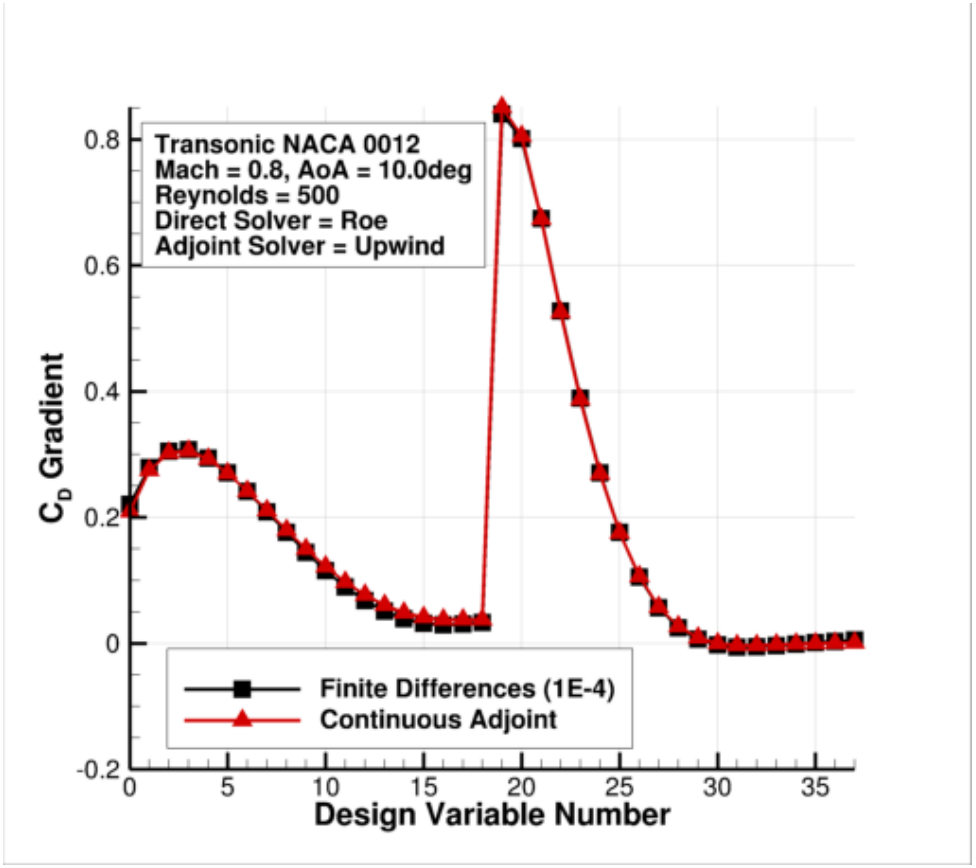

Surface Formulation Assessment (Turbulence)

- While formulations that explicitly treat turbulence (S-A model) have also been developed, what happens if you “freeze” the viscosity?
- With experience, we have found that in many situations (especially compressible, high Reynolds number flows), this is a good option due to its sufficient accuracy and efficiency/lower complexity.

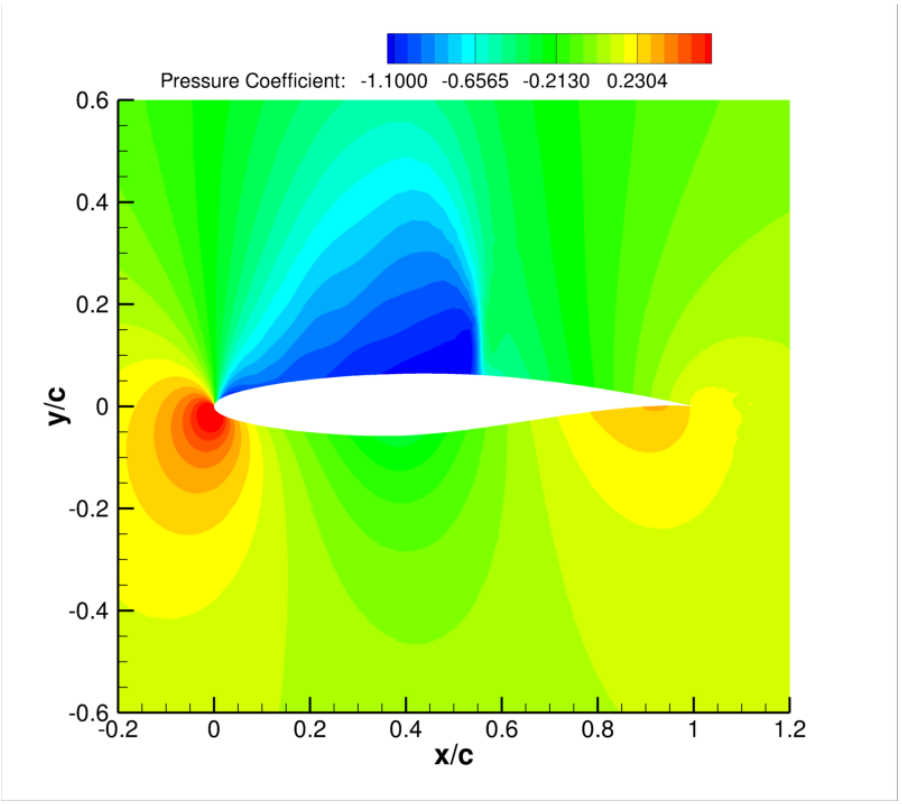

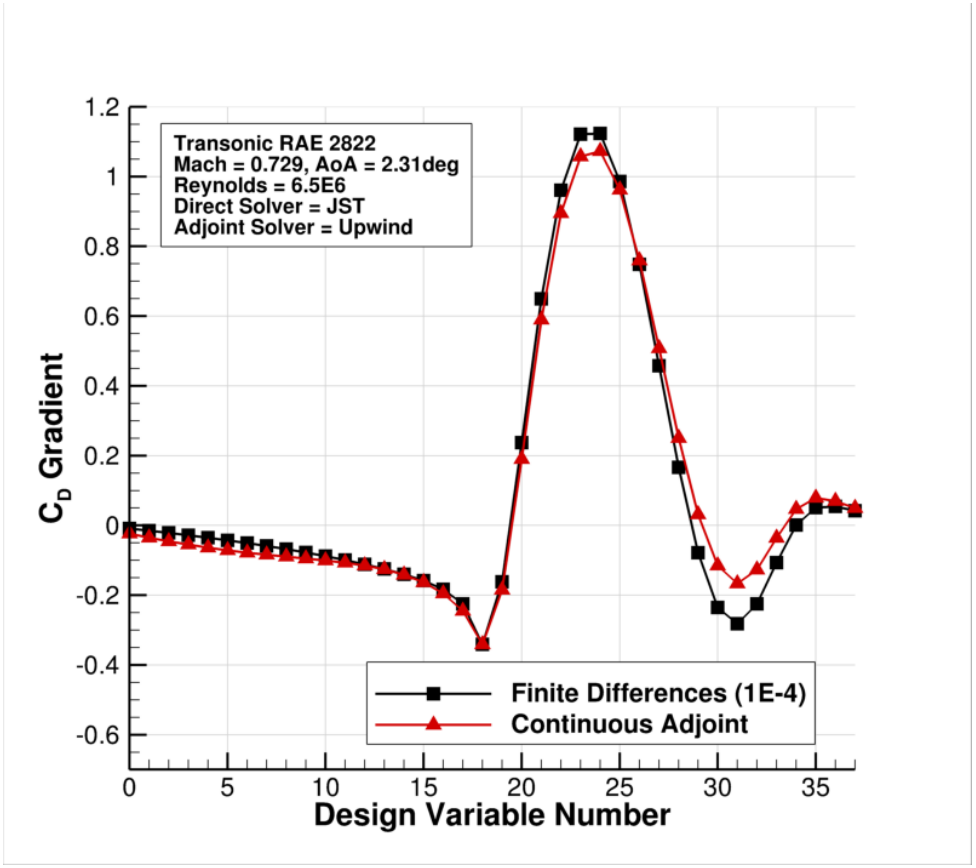

Opportunities in Accuracy / Robustness

- We know that **sharp edges** can sometimes pose a problem for the surface formulation:
  - A number of techniques have been explored for dealing with inaccuracies in the gradient due to these locations, e.g., removal of these points from sensitivity formula.
  - Sharp edges could potentially be treated mathematically, but we have not made this extension.

- **Numerical methods** for the convective terms of the continuous adjoint PDE:
  - These terms are typically treated in a non-conservative fashion with modified centered or upwind schemes.
  - **In our experience, these terms are also the culprit in divergence**, especially near leading edges, sharp edges, or poor quality mesh cells!
  - We are pursuing advanced numerical methods with custom dissipation sensors or upwind limiters for controlling this behavior.

- What about complex problems with flow separation or multiple physics?
  - **Discrete adjoints** might be the best choice here, when numerically exact gradients are needed or the continuous derivations become overly complicated / impossible.
  - Comes at a cost of efficiency in compute and memory, but opportunities remain here for improving their performance as well.

1) Zhou, Y., Albring, T., Gauger, N., Economon, T. D., Palacios, F., and Alonso, J. J., "A Discrete Adjoint Framework for Unsteady Aerodynamic and Aeroacoustic Optimization," AIAA Aviation Forum 2015, Dallas, TX.
2) Additional work recently submitted to AIAA SciTech 2016.

Design Using the Surface Formulation

- With this formulation it is possible to optimize airfoils or isolated wings using the design tool as a black box!
- To do so, we choose suitable methods for geometry parameterization (design variables), mesh deformation, flow solution (PDE), and adjoint solution (PDE) and let a gradient-based optimizer drive the bus.

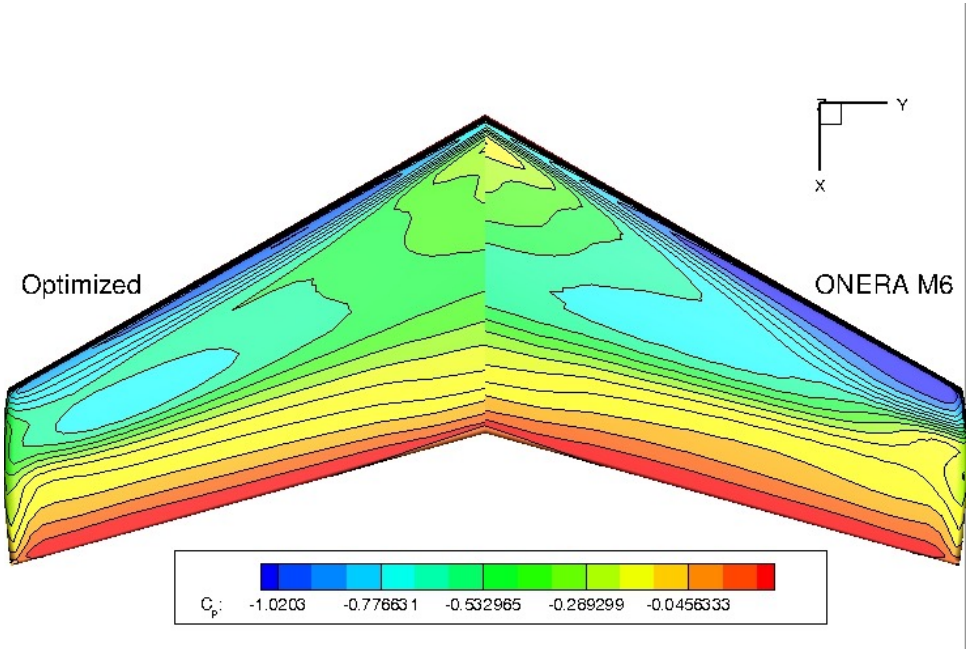

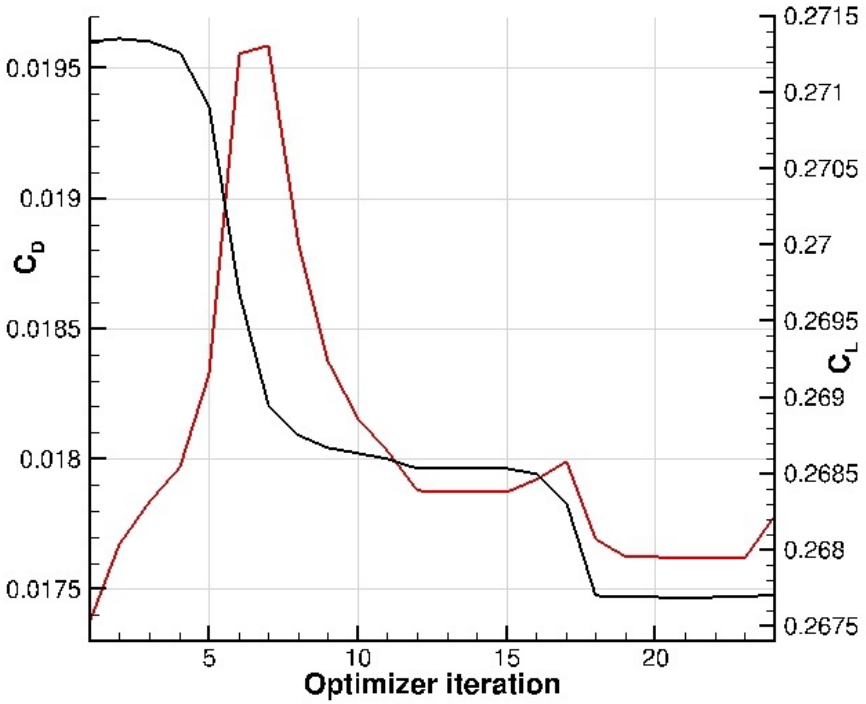

22 drag counts less, same C_L and satisfying geometrical constraints

Redesign of the NASA Common Research Model

Goal: Redesign the wing and tail of the CRM for reduced drag in a number of scenarios that include realistic lift and pitching moment (trim) aerodynamic constraints along with wing thickness constraints using hundreds of design variables (FFD).

Palacios, F., Economon, T. D., Alonso, J. J., "Large-scale aircraft design using SU2," AIAA Paper 2015-1946, 53rd AIAA Aerospace Sciences Meeting, AIAA SciTech, Kissimmee, FL, January, 2015.

Some Highlights:

- Meaningful designs reached for all cases.
- Optimizer had a noticeable impact on results (scaling).
- Starting from different initial angles of attack allowed exploration of local minima.
- Constraints should be added incrementally in order to understand the problem…
- **The role of the designer remains critical in guiding the process!**

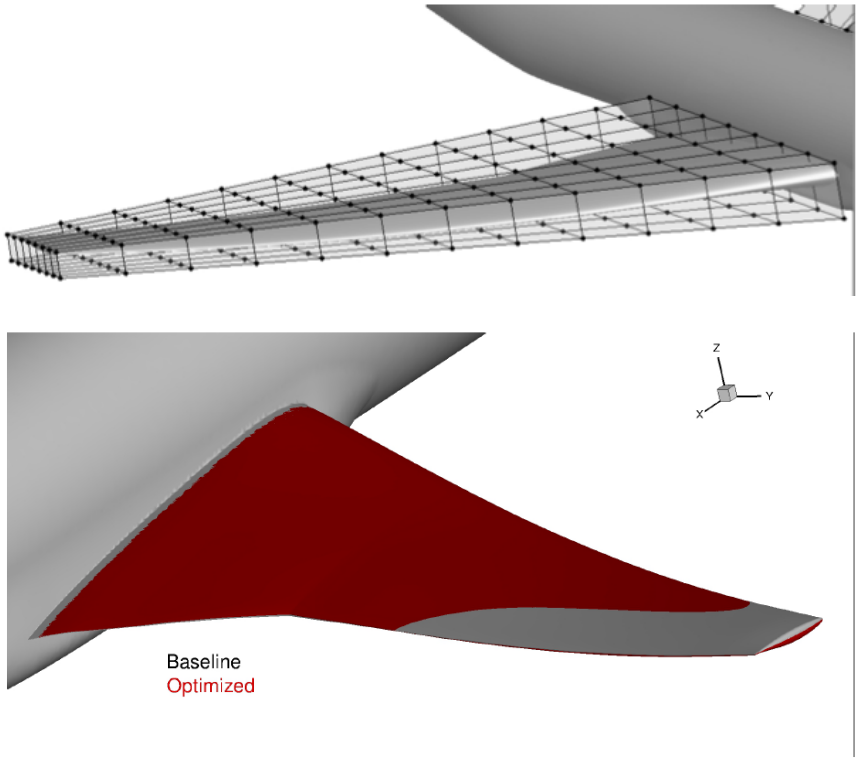

What about performance?
Intel Parallel Computing Center at Stanford

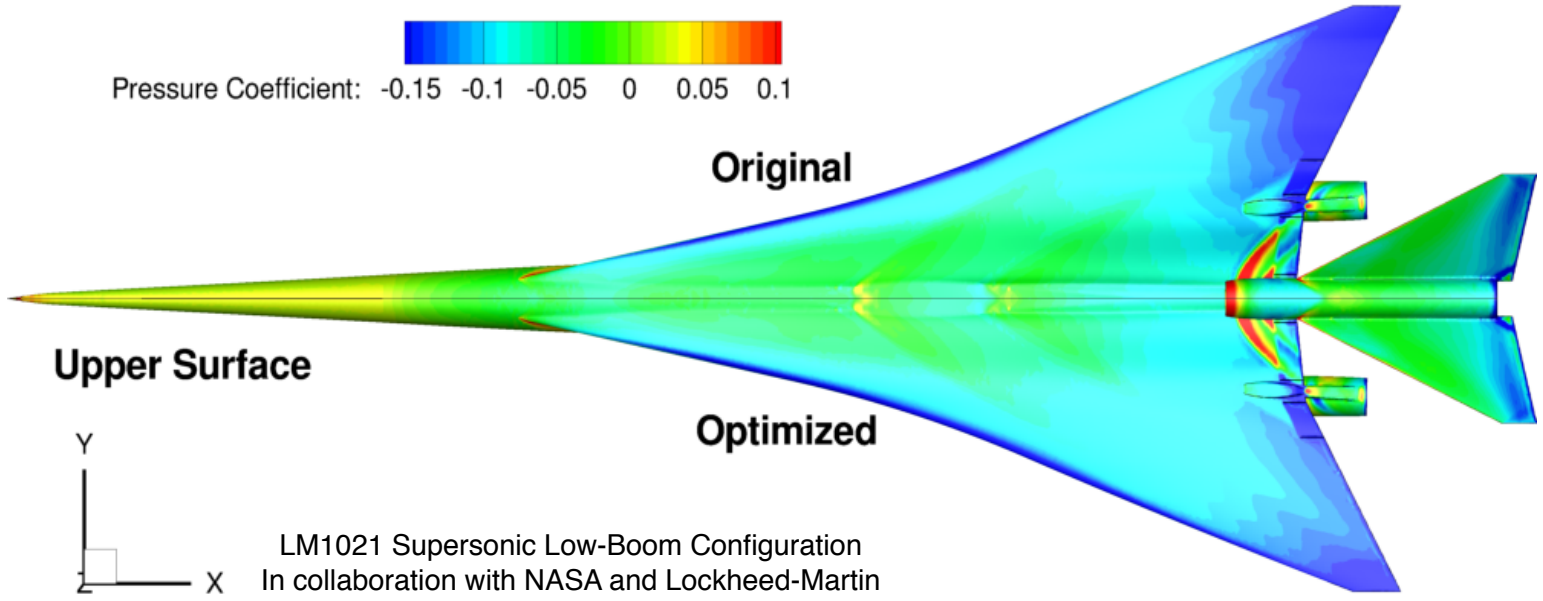

LM1021 Supersonic Low-Boom Configuration
In collaboration with NASA and Lockheed-Martin

- Continuous adjoint in SU2 has no memory or time overhead compared with the flow solver – many solver components are reused for the adjoint. **High-performance code optimizations can be immediately leveraged by the adjoint solver.**
- Potential for **everyday design and optimization** of complex engineering applications on a desktop workstation: revolutionize engineering practice.
- Provide **examples of a scalable, high-performance, open-source solutions** to engineers and computational scientists around the world: exploit significant growth in available computing resources.

Single-Node Performance Optimizations

Key Idea: as modern architectures move toward many-core with more hardware threads and additional vector lanes, **one must employ parallelism at multiple granularities simultaneously to extract performance**.

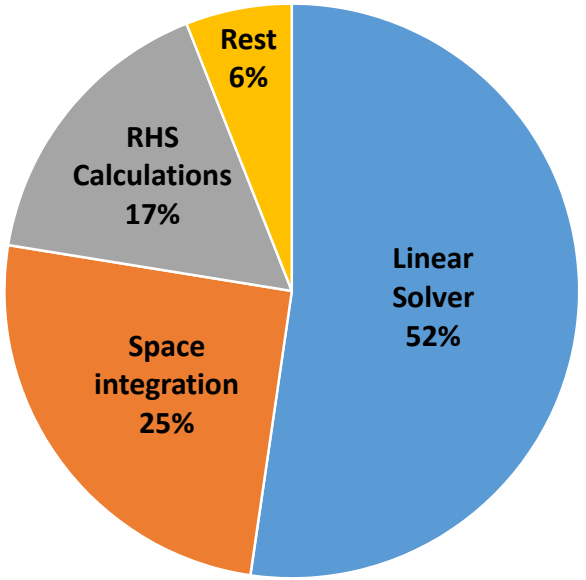

Single-node optimizations on both Xeon, Xeon Phi, for **Implicit RANS**

- ONERA M6 test case was used for these experiments
- Hot spots were identified with custom profiling/scaling tests

➢**Threading: Hybrid MPI+OpenMP**

- Improved cache usage, finer control with working-sets, more flexible work partitioning

➢**Efficient memory (cache) utilization** - improve spatial/temporal locality

- RCM with vertex ordering, Optimal data-structures: AoS → SoA

➢**Load-balancing** - optimal scheduling and work distribution

➢**Vectorization** - effectively utilize the available vector compute units

- 8-wide on Xeon, 16-wide on Xeon Phi
- Beneficial only for the compute intensive kernels – Edge loops, vectorized across edges using SIMD-enabled “Elemental Functions”

Speedups from various optimizations

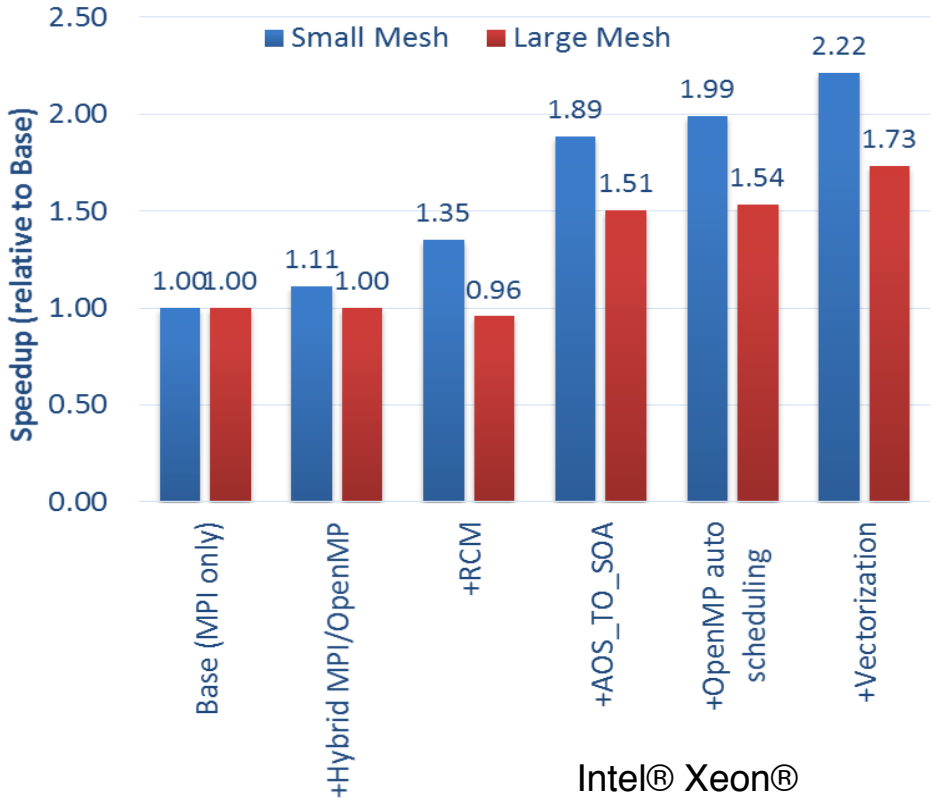

Intel® Xeon®

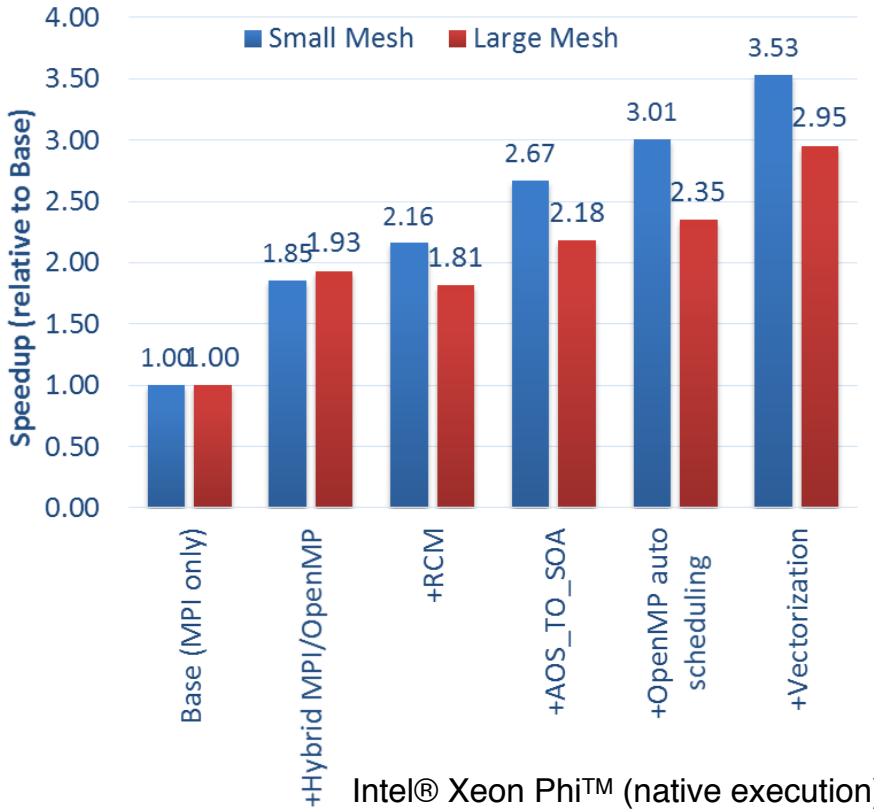

Intel® Xeon Phi™ (native execution)

- The base code is run on 48 MPI ranks on a single Xeon node and 240 ranks on a single Xeon Phi ("flat MPI")
- All others are in hybrid MPI+OpenMP mode: 4 MPI ranks x 12 OMP threads / rank (Xeon) and 4 MPI ranks x 60 OMP threads / rank (Xeon Phi).
- Results here are for explicit RK, inviscid ONERA M6: Small Mesh (94,493 points), Large Mesh (818,921 points) with 100 non-linear (outer) iterations

Scalable Algorithms

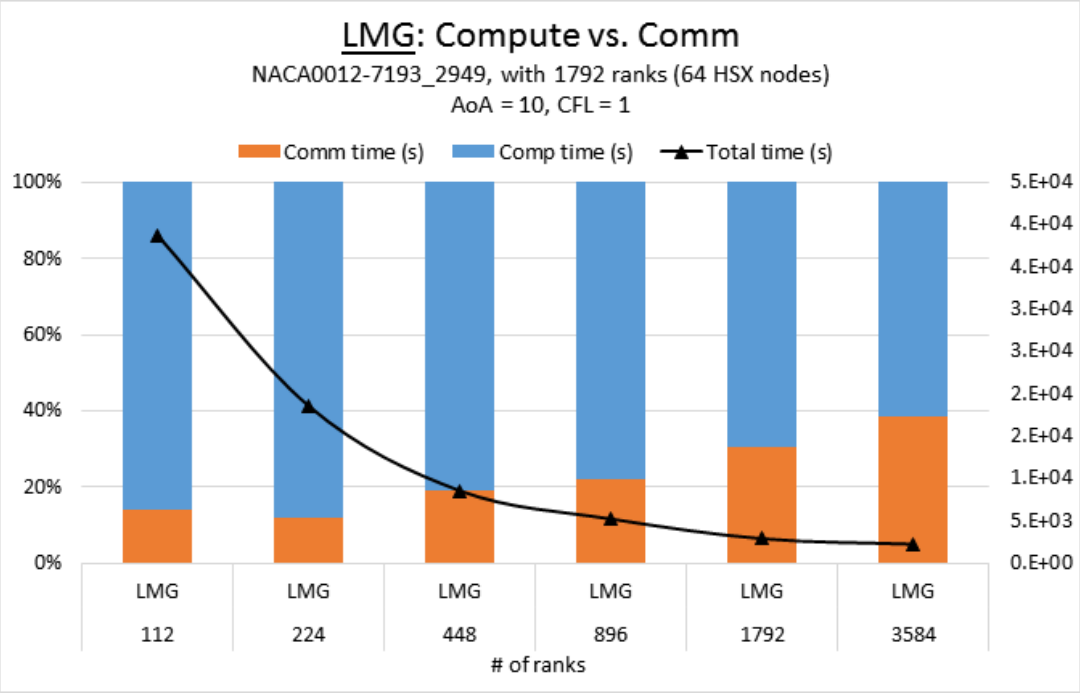

- Choice of algorithm has a major impact on scalability, **linear solvers in particular.**
- We have been exploring the scalability of various algorithms for **implicit** RANS calculations: preconditioners, linear solvers, smoothers, etc.
- Recent head-to-head comparisons show that a geometric linear multigrid solver has better scalability potential than traditional preconditioned Krylov-based methods.

1) Economon, T. D. , Palacios, F., Alonso, J. J., Bansal, G., Mudigere, D., Deshpande, A., Heinecke, A., Smelyanskiy, M., "Towards High-Performance Optimizations of the Unstructured Open-Source SU2 Suite," AIAA Paper 2015-1949, AIAA Infotech at Aerospace, AIAA SciTech, Kissimmee, FL, January, 2015.
2) Additional work recently submitted to Supercomputing 2015 and AIAA SciTech 2016.

Future Directions in Adjoints Conclusion

- This work has successfully demonstrated the use of the Navier-Stokes adjoint methodology in the SU2 suite for large-scale aerodynamic shape design with a realistic aircraft configuration, but…

- **Opportunities exist** for improvements in accuracy/robustness:
  - Sharp edges
  - Advanced numerical methods
  - Discrete Adjoints

- Given the current optimizer technology, the **intuition, experience, and decision-making ability of the designer remain critical in this type of complex, industrial design process**.

- Research in **high-performance computing** for CFD can directly improve continuous adjoints
  - Trend: leverage parallelism at multiple granularities simultaneously
  - Exploring scalable algorithms is critical for large-scale, industry-relevant applications